UC Irvine
UC Irvine Electronic Theses and Dissertations

Title
NeuroFlow: Deciphering Boiling Patterns from Neuromorphic Events

Permalink
https://escholarship.org/uc/item/6ck7w36g

Author
Arani, Srikar

Publication Date
2024

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA,
IRVINE

NeuroFlow: Deciphering Boiling Patterns from Neuromorphic Events

THESIS

submitted in partial satisfaction of the requirements
for the degree of

MASTER OF SCIENCE

in Electrical and Computer Engineering

by

Srikar Arani

Thesis Committee:
Associate Professor Aparna Chandramowlishwaran, Chair
Associate Professor Yoonjin Won
Assistant Professor Hyoukjun Kwon

2024

© 2024 Srikar Arani

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This thesis would not be possible without the support and guidance of several people who, in one way or another, made this work possible. It has been a pleasure to work with everyone who made this thesis a reality.

I would like to thank my advisors Associate Professors Aparna Chandramowlishwaran and Yoonjin Won for their endless help in putting me in a position to finish my thesis and for instilling in me a love for research and scholarship. Their work and excitement with respect to academics stands unparalleled. I would also like to thank my committee member, Assistant Professor Hyoukjun Kwon, whose work and teaching helped bring this work about. Without their guidance and hard work this thesis would not be possible.

The conclusions of this research would not be so enriching without the work of my collaborators. I would like to thank Youngjoon Suh for his wealth of knowledge in this field, and his countless hours explaining the context motivating this work. Further thanks are extended to Sanghyeon Chang for his dedication and immeasurable work in running experiments for this research.

A great many people provided a wonderful source of distraction whilst completing my thesis. A huge thank you is extended to my amazing labmates Alex, Arthur, Neel, Saba, Shakeel, Sebastian, and Zitong. Their jokes, stories, and company, made Irvine a wonderful place to live and work.

The most significant acknowledgement is reserved for my family, whose love and support have enabled me to finish my thesis and accomplish my goals.

VITA

Srikar Arani

EDUCATION

Masters of Science in Computer Engineering **2024**
University of California, Irvine *Irvine, CA*
Thesis: NeuroFlow: Deciphering Boiling Patterns from Neuromorphic Events

Bachelor of Science in Computer Engineering **2022**
University of California, Santa Barbara *Santa Barbara, CA*

WORK EXPERIENCE

Applied Space Researcher **2023**
Johns Hopkins University Applied Physics Lab *Laurel, MD*

- Intern for JHU/APL, in Applied Space Research Group (Top Secret Clearance Group)
- Designed and tested experiments focused on hypervelocity kinetic impact simulations
- Bridged gap between physicists and high performance computing in running dozens of hydrocode simulations
- Aided researchers in gaining a deeper insight into the effect of various topologies on asteroidal bodies
- Developed software suite for analyzing physical properties of large hydrocode datasets

Quality Assurance Engineer **2020 – 2022**
Appfolio Inc. *Goleta, CA*

- Intern for Appfolio, took on the role of the only QA for a team of 8 focused on product security
- Designed and tested in an agile environment and helped push hundreds of projects and stories to users
- Converted internship to part time job while working as a student
- Returned following year to continue working on product security, while also refactoring hundreds of thousands of lines of tech debt to improve system software architecture

RESEARCH EXPERIENCE

Master's Student Researcher **2022-2024**
University of California, Irvine *Irvine, CA*

ABSTRACT OF THE THESIS

NeuroFlow: Deciphering Boiling Patterns from Neuromorphic Events

By

Srikar Arani

Master of Science in Electrical and Computer Engineering

University of California, Irvine, 2024

Associate Professor Aparna Chandramowlishwaran, Chair

Efficient thermal system design is pivotal in cooling applications. The complex nature of boiling phenomena combined with the emergence of diverse flow patterns or *regimes* pose challenges. Despite advances in classification methods, machine learning (ML) models struggle to discern flow regimes from optical data obtained via high-speed cameras.

This thesis presents a compelling case for learning from neuromorphic event representations of flow boiling, which offers insights into hidden properties of flow patterns. We delve into the application of diverse ML algorithms on event data for the task of flow regime classification. Through an analysis of convolutional neural networks, long short-term memory models, and event-based spiking neural networks, we highlight the superior accuracy and performance achieved by event data classification, as well as the unique insights provided by a novel Fourier-based classification approach. We evaluate the different ML algorithms on two event datasets — a video-to-event emulated dataset covering five flow boiling regimes, as well as a dataset combining the emulated data with a neuromorphic event camera dataset encompassing two regimes, for a total of seven flow boiling regimes. This research not only advances our understanding of flow boiling, but also showcases the potential of leveraging new data representations.

Chapter 1

Introduction

The design of efficient thermal management systems represent a crucial aspect to many industrial processes. Core to these structures are flow boiling systems — systems in which fluid is circulated over a heating surface in an effort to dissipate heat, as well as regulate the temperature of a system [1]. Two-Phase flow boiling systems — so-called due to their consisting of a liquid and a vapor phase in the working fluid — are often highly complex and subject to myriad factors. Aspects such as channel geometry, working fluid density, flow rates, and surface tension, all can contribute to experimentation and analysis becoming highly variable [2, 3]. Regardless, as these systems operate, the working fluid absorbs heat from the heating surface, continuously saturating with heat energy, while simultaneously exhibiting various distinct flow patterns or flow regimes. These flow regimes serve as markers for unique heat transfer and dissipation mechanics, making them critical for understanding, designing, and optimizing efficient heat sinks and thermal management systems [4].

Advancements in machine learning (ML) are revolutionizing the identification of these flow regimes; what were once subjective verbal descriptions of visual scenes can now become more objective and data-driven, promising greater consistency and deeper local insight. Efforts

to do flow regime classification traditionally have involved analyzing signals from pressure and temperature sensors, to train ML models capable of recognizing regimes based on the unique signatures each one presents. However, a potential limitation of relying on sensor data lies in its limited capacity for capturing local details within the predicted data, potentially overlooking critical local phenomena that influence effective heat transfer [5, 6]. To bridge this gap, researchers have proposed using image-based models that capitalize on the rich, localized information provided by high-speed optical cameras, all while being less subject to the aforementioned experimental variability [7, 8].

Despite the advantages that conventional image-based models offer in categorizing flow regimes, they are accompanied by several limitations that affect their practicality and efficiency [9]. The operation of high-speed cameras, for instance, requires significant power, which may render them less suitable for applications where energy conservation is crucial or for monitoring over extended periods. Additionally, the visual distinctions between various flow regimes may be hard to discern, even by a trained observer. A traditional image-based model may require both highly detailed and manicured information, as well as a deep or computationally expensive network for accurate classification — luxuries that may be harder to come by in industrial settings where accurate, low-cost, or real-time information may be vital.

Transitioning from the constraints posed by conventional high-speed cameras, neuromorphic event cameras present a novel way of collecting spiking event data from neuromorphic hardware, possibly presenting new techniques for regime classification, as well as providing low-power, low-latency networks. Distinguished by their motion sensitivity, selective data capture that reduces irrelevant background noise, and markedly lower power consumption and memory requirements, they provide a solution for many issues in traditional high-speed cameras [10]. Neuromorphic cameras excel in capturing changes in a scene with unparalleled temporal resolution, focusing solely on pixel-level brightness variations. This approach not

only filters out static information, enhancing the relevance of captured data but also significantly diminishes power usage — from watts down to milliwatts, offering a drastic reduction compared to traditional methods [11]. Moreover, the sparse, event-driven data generated by these cameras necessitates substantially less memory, often just a fraction of what is required for storing high-speed camera footage. This efficiency in both power and storage, combined with their exceptional sensitivity to dynamic changes, makes neuromorphic event cameras a compelling choice for the classification and analysis of flow regimes, promising a leap forward in thermal management research and applications with their high-resolution, energy-efficient capabilities.

In this study, we hypothesize that neuromorphic event data provides a better input for flow regime classification, as opposed to traditional high-speed cameras. In this thesis, we will refer to the neuromorphic event data as Event Data, and the high-speed camera data as Optical Data. In determining the efficacy of event data representations for flow boiling classification, it makes sense to highlight, not just accuracy, but performance via latency and power consumption and even model size. Given the sparse nature of event data, leveraging these metrics paints a more complete picture of its strengths and weaknesses. To do so, we compare various networks, each with their own strengths and weaknesses. Convolutional Neural Networks (CNNs) are more traditional machine learning methods, but they may not best take advantage of the event representation for power, latency, or model size. Long Short-Term Memory (LSTM) networks can make use of temporal trends, however, they struggle in highlighting spatial information. Meanwhile Spiking Neural Networks (SNNs) are purpose built for event data, designed to be very low power, but may not be able to replicate advanced or complicated models. Finally, we propose a novel Fourier-based approach that maximizes the unique signal properties of event data to create a fast and efficient approach to classification.

We evaluate the different ML models on two event datasets, one created via emulation to

leverage existing experiments, and the other derived from real world experimentation as captured by an event camera. Traditionally, the primary mechanism by which the majority of event datasets are generated makes use of emulators or recording existing datasets with event cameras. While emulated datasets are very useful, they may not be as representative of real world conditions, as opposed to true event camera captures of experiments, due primarily to preexisting frame rate of the optical data, as opposed to the high temporal resolution of true event cameras [12]. Regardless, we find that event data representations provide a new and powerful way of approaching the crucial problem of flow regime classification.

Chapter 2

Related Work

The area of neuromorphic event data and event-based computing is relatively new. Even still the field is rife with innovation. Additionally, flow boiling classification is another field that similarly provides many existing frameworks to build from.

2.1 Flow Regime Classification

Flow Regime Classification has a long history, and there exist many approaches for this problem. Many older methods made use of sensor arrays to gather physical information about the inside of a channel, though this technique has slowly been phased out for optical data. Huo et. al. present a classification method for determining heat transfer coefficients using the physical characteristics of the vapor in seven regimes [13]. This presents an interesting approach to extracting physical characteristics of different flow regimes, however, this paper does not present a machine learning approach for real time classification. Additionally, the lack of optical data may make this style of classification too specific to the environment in which it was trained.

Hernández et. al. [14] provides an interesting look into more traditional and older methods for regime classification. These methods make use of large arrays of sensors, and highly susceptible to the varying conditions of the experiment. Additionally, their paper works on a much smaller subset of flow regimes, failing to reach Annular regimes, a major regime in working with high pressure systems near Critical Heat Flux (CHF). The paper additionally reports unremarkable accuracy, combined with the smaller experimental matrix, and more convoluted measuring systems make this an interesting look into traditional classification methods, and an area in which we can greatly improve.

In their paper, Ooi et. al. present a machine learning approach to flow regime classification [15]. They present multiple methods that require large sensor arrays in the same vein as the paper by Hernandez et. al. They identify three regimes with this data, and achieve an accuracy of 88.4%. Additionally, as other authors highlight, unique experimental design or channel environments make a sensor approach too specific to a single rig. Using optical data, and classifying more regimes, at a higher accuracy, are all areas that can be improved to better understand flow boiling patterns.

Hobold and Silva present a newer technique in tackling this problem [16]. They utilize Principal Component Analysis (PCA) for a huge dimensionality reduction and are able to achieve accuracies above 90% in their classification. However, they also classify only 3 regimes, and, while they also utilize optical data, their accuracy has room for improvement. However, their heuristic approach with PCA might prove to be an interesting addition to further work in refining heuristic event data approaches such as our Fourier-based classifier.

2.2 Event Data Models

With regards to event data models, Bi et. al. [12] provides an interesting approach to designing a graph based CNN for for object classification using event data. This work is very similar to our own in that it presents a unique network to build around event data. This model and their accompanying dataset, however, is designed for American Sign Language (ASL) sign classification. Additionally, it fails to compare the event data representations to more traditional optical networks to showcase if and how a traditional network may fare at the same task. While it is a motivating work in event-based models, we hope to expand on the idea with direct comparison to current state of the art traditional models on the same task.

Sabater, Montesano, and Murillo designed the Event Transformer [17]. This network is a large event and multi-modal network that performed favorably against other event models in event gesture classification. This network, while impressive, also fails to compare an event data representation to an equivalent optical data representation. Further, like many event datasets, these models are not designed for observing and classifying physical processes. The unique issues in physical data are further expanded upon in this thesis.

Further, Gamage et. al. propose using event data for detecting defects in civil infrastructure [18]. They compare traditional CNNs against event-based SNNs, and highlight the increased accuracy and speed of event-based models for classification. However, they only compare optical CNN's against event SNNs, and fail to interrogate the best architectures for event data. Additionally, they do not compare state of the art CNNs against SNNs. We expand on this work by both comparing various event-based architectures, as well as providing traditional optical CNNs state of the art networks.

On the other hand, Liu et. al. have used event-based data to design a fast and efficient data representation and associated network for event-data classification [19]. This work compares

their methods against many other event-based networks and highlights the superior accuracy of their own. However, like most other event data research, it is evaluated on large benchmark datasets, many of which are emulated versions of existing datasets. Additionally, while their work provides accurate, power and time efficient classification, it is a large network designed to compete against other state of the art networks, and is poorly suited to work on edge devices in labs or experimental setups. Xie et. at. also design a large network to leverage event data, and much like Liu et. al, they compare their event network against multiple other event representations [20]. However, they also design a network designed to compete against state of the art networks. These networks also are not compared to equivalent optical networks. Thus, while these networks are highly accurate, they poorly articulate why an event representation would be preferable to an optical dataset. We expand on this work by highlighting a unique use case in which event data provides a unique insight into the problem domain.

2.3 Event Data Models for Flow Boiling

Finally, a paper by Lu et. al. presents a very close analogue to the problem being addressed [9]. They present the case for an event-based approach to determining whether or not the Critical Heat Flux condition has been reached in pool boiling experiments. Pool boiling is a similar field to that of flow boiling and this paper provides motivating results, highlighting the value of event data for speed, while retaining accuracy. We expand on this paper by tackling a problem wherein traditional methods have thus far struggled much more, and extend the problem from a binary classification to a multi-class regime classification problem. Additionally, we present multiple machine learning models, as well as the case for a more analytical heuristic approach to leverage and better understand the signal and physical components that make up these regimes.

Chapter 3

Flow Regime Dataset

We create two event datasets for the purpose of flow boiling regime classification. The first, consists of data captured by a Prophesee neuromorphic event camera aimed at an experimental rig. The specifics of the experimental setup are discussed in more detail in Appendix A. The second, is developed by processing experimental data through an emulator in order to generate an approximate event stream. The specifics of the video to event (V2E) conversion process are discussed further in Appendix B. The former consists of data comprising two distinct flow regime classes, while the latter comprises of five distinct flow regime classes — encompassing seven total classes as seen in Figure 3.1. For comparative analysis, an optical dataset based on the latter is created to present a baseline. In this section we describe both the observed classes, as well as the datasets created from those classes.

3.1 Flow Regime Classes

Flow regimes each exhibit unique flow properties, but there exist many different regimes and sub-regimes, and some experiments may better highlight some regimes as opposed to others.

In this thesis, we restrict our focus to seven regimes — *Annular, Slug, Stratified Wavy, Stratified Smooth, Bubbly, Elongated Bubbly, and Unstable*. These regimes serve to highlight the full range of a flow boiling system, from the beginning of boiling, until approaching Critical Heat Flux — the point in which the heat dissipation process fails. The different regimes, and their representation in both optical data and event data, are highlighted in Figure 3.1.

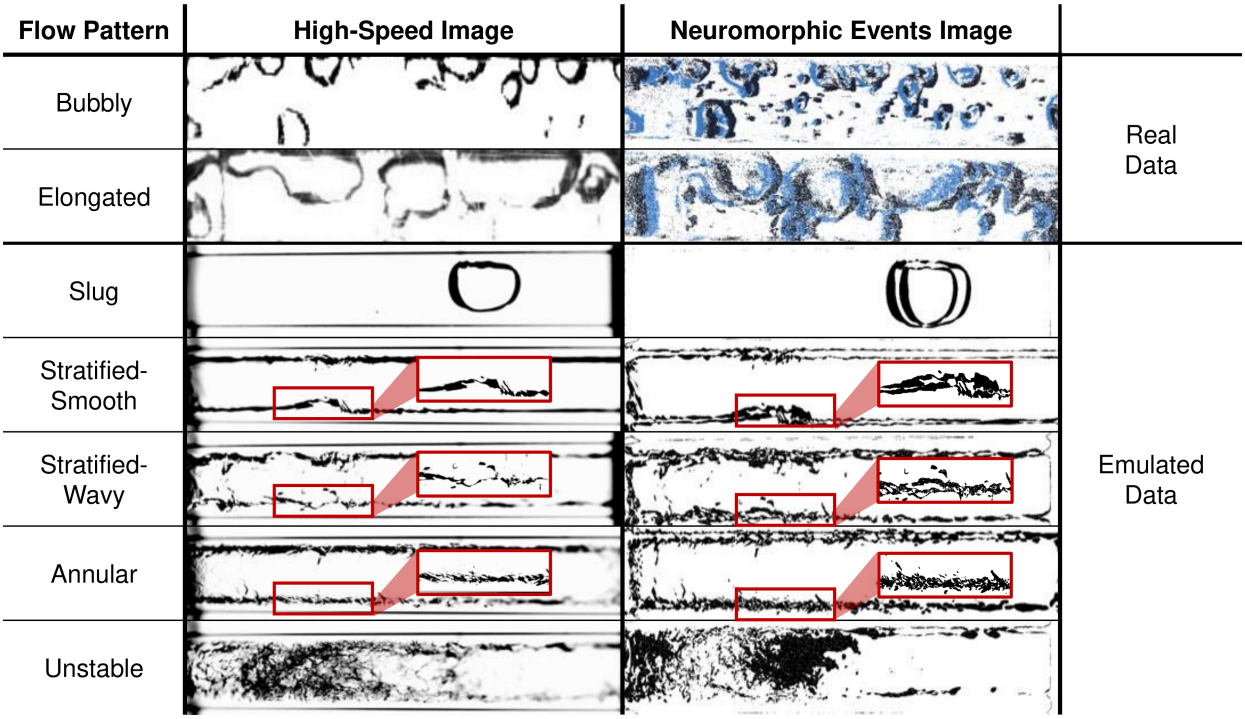

Figure 3.1: **Visualization of the Optical and Event datasets.** Both columns represent the same point in time. Cutouts highlight the similarity in Stratified Smooth, Stratified Wavy, and Annular regimes with respect to high-speed optical images. The differences in these regimes are better highlighted in the neuromorphic event data representation.

3.1.1 The Annular Stratified Problem

One factor that makes effective classification particularly difficult is the similarity in the flow patterns of the *stratified wavy*, *stratified smooth*, and *annular* regimes. This similarity contributes to the difficulty in using optical data from high-speed cameras, however, as can

be seen in Figure 3.1, event representations better highlight the fine-grained differences in the regimes. Note the narrow event bands in the Stratified Smooth regime, which, are due to laminar flow with well defined vapor regions in this regime. However, as the pressure changes, the boundary between the vapor and liquid becomes more turbulent. This results in more clustering of events around the boundary in the Stratified Wavy Regime. By the time the system reaches the Annular regime, the liquid becomes akin to a thin film around a large vapor region. The thin boundary between the liquid and vapor leads to an increase in the frequency of events and more smaller clusters of events around this boundary. These differences are highlighted above in Figure 3.1.

3.2 Emulated Data

In the realm of event based networks, many datasets have been created from the existing datasets. Event datasets such as MNIST-DVS [21], CIFAR10-DVS [22], N-Caltech101 [23], and N-MNIST [23] were converted from the datasets MNIST, CIFAR10 and Caltech101 by displaying the original datasets on a monitor, one image at a time. In the case of N-MNIST and N-Caltech101, an event camera is then moved slightly back and forth, and in the case of MNIST-DVS and CIFAR10-DVS, an event camera is fixed and the displayed image is moved slightly. While this method works fairly well, emulators that can recreate this effect in software have started to present another approach to this issue of creating data for event-based models.

We create a dataset, comprising of five regimes, *Annular, Slug, Stratified Wavy, Stratified Smooth, and Unstable*. While emulated data is not without problems, it allows for leveraging existing experiments in new models. This data is converted from existing experimental data captured by a high-speed camera. Fifteen experiments, each defined by a separate heat and mass flux, are passed through an emulator to generate approximate event streams,

before being sorted for regime classification (Appendix B). These event streams are then reconstructed into *Event Frames* by accumulating 50,000 events, and creating an image, taking the x and y coordinates of each of the events, and marking a pixel at that point (Figure 3.1). This data is compiled in Table 3.1. The specific breakdown of experiments and their emulated frames is compiled in Appendix C.

Dataset Class Breakdown			
Class	Optical Frames	Event Frames	Resolution
Stratified Wavy	1,555	10,203	1024×768
Stratified Smooth	688	3,721	
Annular	2,402	18,032	
Slug	878	6,081	
Unstable	750	2,471	
Bubbly	N/A	8,759	1280×720
Elongated Bubbly	N/A	1,235	

Table 3.1: **Class breakdown for both Optical and Event Data.** The classes are highly unbalanced, contributing to the challenge in accurate flow pattern classification.

3.3 Event Data

We additionally create a second dataset, provided by a real event camera. This dataset comprises two regimes, *Bubbly, and Elongated Bubbly.* These regimes are more common very early in the boiling process, as opposed to the aforementioned emulated regimes. These datasets differ both in their data acquisition techniques, but also in their experimental setup and goals. While both datasets highlight unique flow boiling regimes, they observe slightly different dynamics. The higher temporal resolution of the real event camera allows for better observation of the bubble nucleation process, observed in the earlier flow patterns such as Bubble and Elongated Bubbly. The goal of using both emulated data, as well as real event data, is the design of a model that can accurately capture the full range of boiling experiments, as well as be generalizable in a variety of conditions. The experimental

setup and event capturing process is further expanded upon in Appendix A, and the Event Frame breakdown is highlighted above in Table 3.1. An important note on the data that can be seen in the table, is the highly unbalanced nature of the dataset. This both is a symptom of the difficulty of the problem, and a cause. Obtaining good data in each of these classes is difficult, due in part to the somewhat transient nature of some of these regimes — in particular Unstable flow patterns often occur in the midst of other patterns such as Annular and Stratified Wavy regimes. This contributes to a lack of data in experiments for classification, and is an area in which an accurate flow regime classifier may be used in. The specific breakdown of experiments and their emulated frames is compiled in Appendix C. Additionally, the size of the event dataset is substantially larger than the optical dataset. The difference in dataset size may present a confounding variable for directly comparing the accuracy of each dataset. This is explored in the Ablation Chapter.

Chapter 4

Classification Approaches

To gain a deeper understanding of the efficacy of event representations, we compare multiple machine learning approaches, such as different CNN architectures, an LSTM, an event-based architecture such as an SNN, and an analytical Fourier-based method. In comparing traditional data captured via high-speed camera and neuromorphic event data — both emulated and real — we will refer to the former as Optical Data, and the latter as Event Data. All hyperparameters are highlighted in Table 4.1.

4.1 Data Preprocessing

An important distinction in how both the optical and event data compare, as well as how each of the ML approaches compare, is the preprocessing required for each model. Optical Data consists of individual frames that are processed as images in the CNN as shown in (a) in Figure 4.1. The Event Data, however, is reconstructed into an image in the CNN, SNN, and Fourier-based analytical method. This whole process is outlined in Figure 4.1, highlighting how the original data is converted into reconstructed frames for CNNs, SNNs,

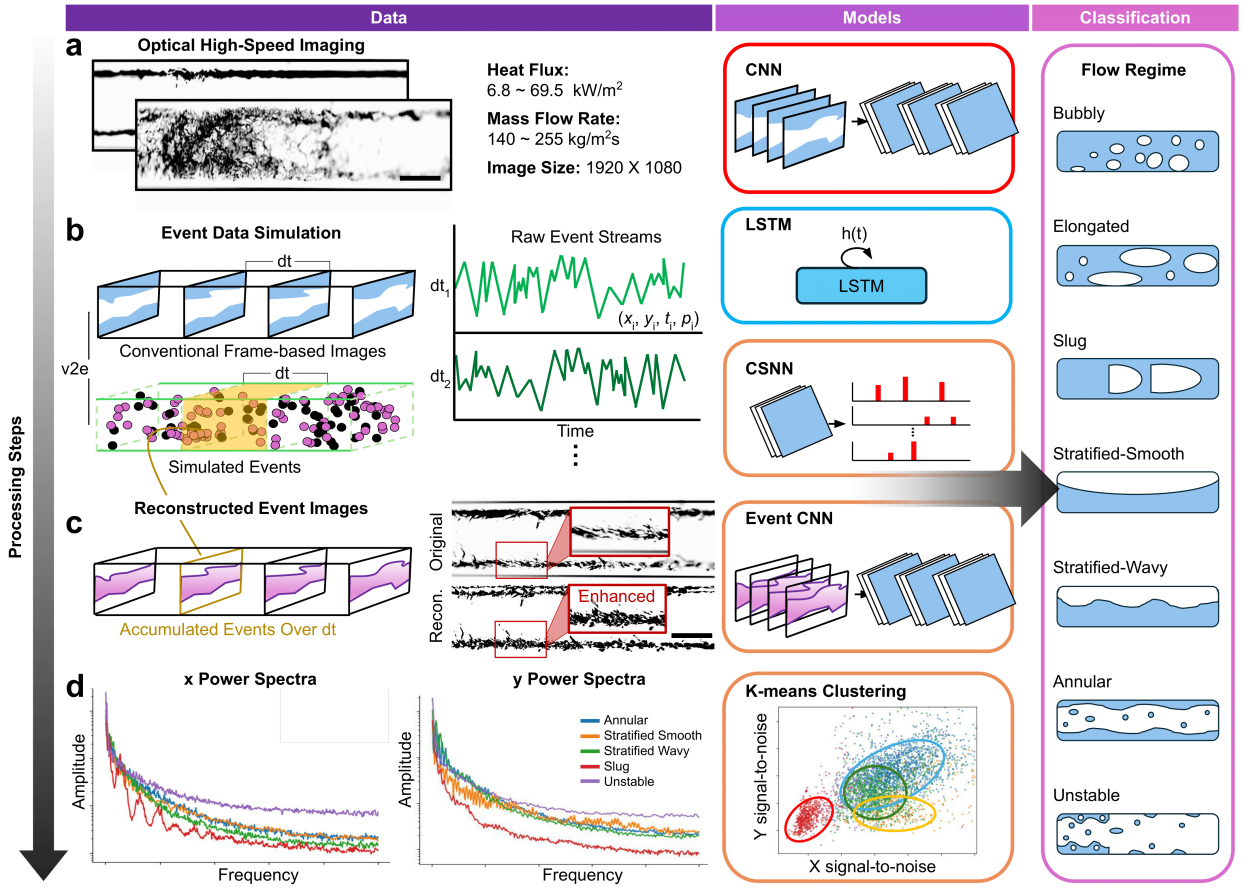

Figure 4.1: **Overarching methodology for flow regime classification approaches.** Both emulated and real event data are processed identically, but differ from optical data. The optical data is only used in the CNN as show in (a). Event data is left as raw events in the case of the LSTM (b). The events are reconstructed into frames for the purposes of event CNNs, Convolutional SNNs, and the Fourier-based analytical method (c). The Fourier-based method, further processes the data into its x and y power spectra (d) before applying a clustering method and finally a k-nearest neighbors classifier.

or the Fourier-based Approach (c), or can be processed as raw events for LSTMs (b). The Fourier-based approach in particular further processes the data into its x and y power spectra for a signal processing approach to the same issue of classification (d). Each ML approach uses a batch size of 128 for training. Additionally, the data precision of the model is left as a 32-bit floating point number. Future work could explore integer quantization for more efficient deployment.

4.2 Convolutional Neural Networks

Convolutional Neural Networks can exploit the spatial structure of an image through the use of convolutional layers, which apply filters (also known as kernels) to the input data, scanning it in a sliding window manner. These filters learn to detect various features or patterns, such as edges, corners, or textures, at different locations in the input. Generally these convolutional layers are followed by non-linear activation functions such as ReLU and downsampling layers such as max pooling. Multiple such convolutional blocks are stacked on top of each other, allowing the network to learn increasingly complex and abstract features. These features can then be used for several downstream tasks such as classification, segmentation and object detection. While event streams are not explicitly designed for CNNs, reconstructing an image from an event stream allows for a powerful point of comparison, as both the optical image dataset and the event dataset can be trained on CNNs.

We experiment with different CNN architectures to determine the best fit for flow regime classification. Two popular CNN architectures are ConvNeXt and AlexNet. These two networks are well established in the realm of classification, ConvNeXt in particular performs favorably against the state of the art Transformer networks at ImageNet classification [24]. AlexNet consists of five convolutional layers, followed by a large fully connected network. While not as popular anymore, AlexNet provides a more simple, but well established model

for comparison. ConvNeXt, however, presents a more transformer-like architecture, composed of multiple ConvNeXt blocks, each of which comprises multiple convolutional layers. Due to the very different nature of the flow boiling classification data, both networks are trained from scratch on the optical and event data. However, we also compare against fine tuning the networks on both event and optical data for additional comparison. Each network downsamples the input data to ImageNet specifications, and the layers of the network are frozen for transfer learning with an Adam optimizer and a learning rate of 1e-3. Further work could include comparing light-weight networks such as MobileNet or EfficientNet.

Additionally, we design a shallower network to address the sparsity of the data. Both the optical data and the event data are downsampled before being passed through a shallow network for classification. Due to the different aspect ratios of the original data, the event data is downsampled to 256×192 px (one quarter of the original dimensions), while the optical data is downsampled to 340×192 px (about one fifth of the original dimensions). The network consists of three convolutional layers, with a maxpooling layer between the first and second. This is followed a single fully connected layer leading to the output. This simpler network is designed for better performance, and lower computational complexity, while also leveraging the relatively sparse nature of the data, and few classes. These networks are also trained with an Adam optimizer and a learning rate of 1e-3.

4.3 Long Short-Term Memory Networks

The same way CNNs can exploit the spatial structure of an image, Long Short-Term Memory (LSTM) networks can exploit the temporal aspects of data. These recurrent networks take into account previous data samples and excel at tasks where sequences of data are meaningful. In boiling, these flow patterns are primarily observed over a period of time. With the exception of unstable regimes, it is unlikely for a flow pattern to quickly transition from slug

to stratified and then to annular, etc. in rapid succession.

We design a shallow LSTM for the event data. Due to the different architecture, the event data is not reconstructed as in the CNN, but left as raw event information. The network input is a vector of 5,000 events, each event composed of 3 values corresponding to the event's x position, y position, and polarity (x,y,p). The network consists of just three recursive layers. This is followed a single fully connected layer leading to the output. This difference in input size was chosen to reduce computational complexity and latency. Like the CNN, this network is trained with the Adam optimizer and a learning rate of 1e-3.

Hyperparameter	Event CNN	Optical CNN	Event SNN	Event LSTM
Epochs	25	35	55	25
Batch Size	128	128	128	128
Optimizer	Adam	Adam	Adam	Adam
Learning Rate	1e-3	1e-3	1e-4	1e-3
SNN β	N/A	N/A	0.95	N/A
SNN Steps	N/A	N/A	15	N/A

Table 4.1: **Generic Hyperparameter Settings for Models**

4.4 Spiking Neural Networks

Given that event data is a focal point of this research, it makes most sense to consider a neuromorphic neural network for a possible classifier architecture such as a Spiking Neural Network. It can be thought of as analogous to some simple ML techniques, however, each of the neurons only represents a binary on or off. This is an extremely power efficient approach, and it may even be highly time efficient. On the other hand, it may be hard to synthesize for real-time embedded systems.

In designing an SNN model for flow-boiling classification, we developed a Convolutional SNN, built with the SNN-Torch framework [25]. Convolutional Spiking Neural Networks have

shown promise as more accurate models than traditional Convolutional Neural Networks at classification tasks [26]. This framework allows for the implementation of a similarly shallow network as the one outlined in Section 4.2, however, leaky integrate and fire neurons are added between layers to create spiking activation. This network provides a possible way to increase accuracy, while reducing computational complexity or power consumption, though this is harder to observe on traditional hardware, than it is on neuromorphic hardware.

The SNN model is trained on the event data, preprocessed and downsampled in the same manner as the CNN. The leaky integrate and fire neurons receive inputs over time and decay until they receive another sufficiently high impulse. The flow boiling regime classification SNN had a decay rate of 0.95. The network is trained with the Adam optimizer and a learning rate of 1e-4. The model consisted of two convolutional layers, each followed by a maxpooling layer. After each of these layers, the signal is passed to a spiking neuron layer. Finally the model consists of a single fully connected layer, followed by a spiking neuron layer, eventually leading to the output.

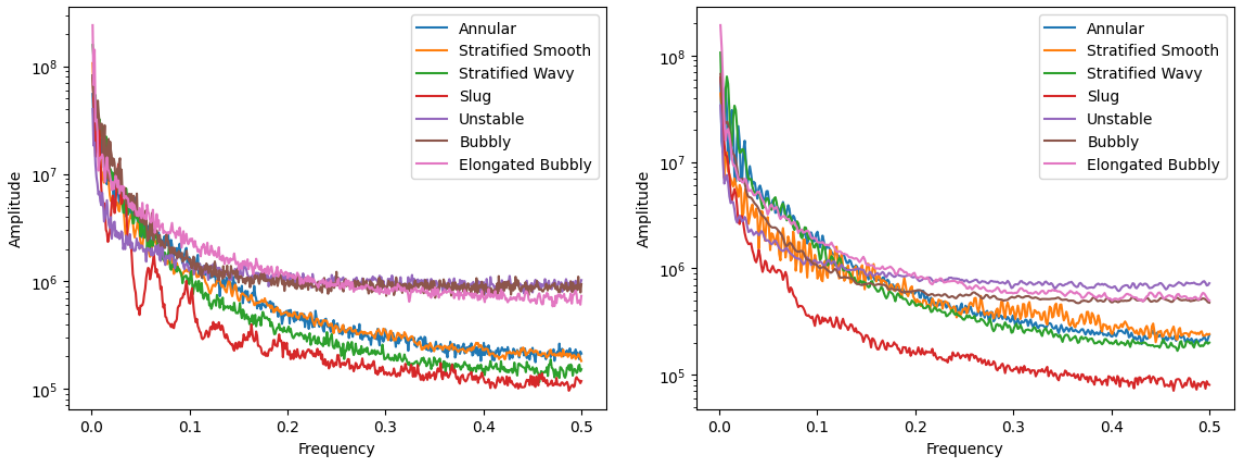

(a) Across x dimension

(b) Across y dimension

Figure 4.2: **Power Spectra for five emulated regimes.** These power spectra highlight representative examples of the x and y power spectra

4.5 Fourier-based Analytical Method

In this section, we propose a novel method for dealing with event data, highlighting the unique signal properties of event data. Given the difficulty of regime classification in the optical representation, especially with regards to the *Annular-Stratified* problem, an efficient event-based classification model should build around the unique benefits of event data. As highlighted earlier, there are unique noise properties in the signals that can be built on — the increased noise in the Annular regime, or the much reduced noise in the Stratified Wavy Regime — for example. Additionally, while some regimes consist of primarily horizontal data such as the three aforementioned problem regimes, others such as Unstable Flow consist of large amount of noise in both the horizontal and vertical dimensions (as can be seen in the power amplitude in Figure 4.2a), or Slug Flow which consists of much less noise and power in the vertical dimension (as can be seen in the power amplitude in Figure 4.2b). These unique properties make event data a great fit for a signal processing approach that can more easily highlight this information. Additionally, given the importance of fast and efficient classification, using a Fourier Transform based approach may be significantly faster than traditional machine learning techniques.

Like the CNN, this method requires reconstructing the event stream into a series of event frames. We then perform a Fast Fourier Transform (FFT) along both the x and y dimensions of the frame, and finally compute the power spectra for the data in its x and y dimensions. The resulting spectra are both highly consistent among their regimes, and slightly different between different regimes. An example of this can be seen in Figure 4.2, comparing the power spectra in the x dimension for each of the five emulated regimes. These signals present an opportunity for exposing more of the unique signal properties of event data. The combination of factors such as signal-to-noise ratio, average power, peak power, and many others, across both the x and y dimensions, lend themselves to clustering algorithms for a more heuristic approach for flow classification. The signal-to-noise ratio, as one example, is particularly

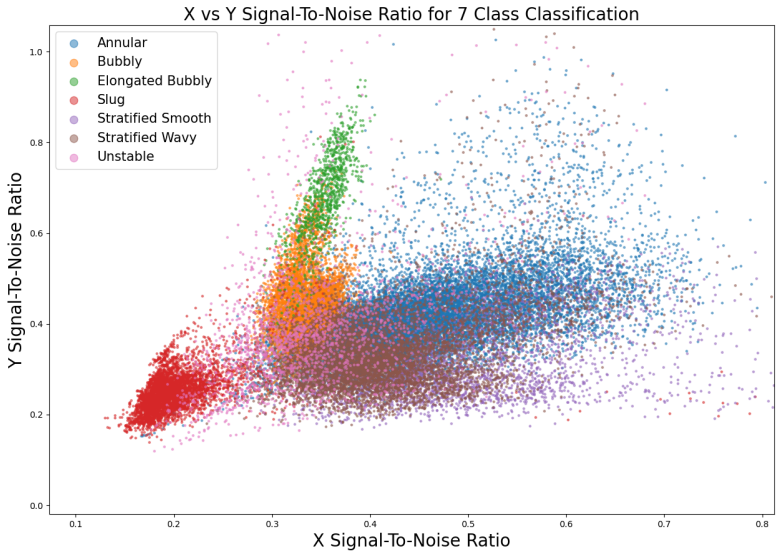

Figure 4.3: **Example of clustering across the x and y signal-to-noise ratios.** The Slug regime, depicted in red, shows dense clustering in this particular example. These are two of many factors that create the clustering necessary for the Fourier-based analytical method.

good at highlighting slug regimes, as opposed to the other regimes as can be seen in Figure 4.3. The combination of these factors across the entire training dataset are then compiled. A k-nearest neighbors algorithm is then applied to this multi-dimensional dataset using 7 nearest neighbors. This additionally exposes some of the more opaque physics aspects of the problem, for example, highlighting the signal-to-noise ratios between regimes exposes the high frequency vapor components of the various regimes.

Chapter 5

Results

Model	Accuracy	Inference Time	Model Size	Model Parameters
Event CNN	**99.04%**	0.15ms	1.90MB	497,413
Optical CNN	92.24%	2.55ms	2.51MB	658,837
SNN	98.92%	0.76ms	**0.94MB**	**247,013**
Fourier	84.68%	**6.44**μs	5.88MB	N/A
Event LSTM	92.54%	1.03ms	1.27MB	332,933

Table 5.1: **Final results comparing 4 different methods for flow regime classification.** Bold terms represent best performing metrics.

In determining the real-time and practical viability of each of these models for flow regime classification, it is important to consider different metrics such as inference time and model size. This is also important for comparing event data representations to traditional high-speed camera representations, as the trade offs inherent in each may provide different choices for different applications and use cases. Given that flow boiling experimentation and observation happen in labs and industrial systems, it is important to not just highlight accuracy, but compare each model's ability to run on any edge device, and in real-time critical conditions. The results discussed here are for the emulated dataset for a direct comparison with the optical dataset. All models were trained and evaluated on an Nvidia A30 GPU.

The event CNN and SNN performed almost identically, to each other in terms of accuracy ($\approx 99\%$), and both are more accurate than the optical CNN ($\approx 92.2\%$). Where they differ, however, is in their inference time and model size. The event CNN is a larger model as seen in Table 5.1, but performs inference in a much shorter time. In systems where latency is an important issue, an event CNN might prove to be a more fitting model. However, in systems where space is a more pressing concern, the SNN provides for a much smaller model (about half the size as the event CNN).

On the other hand, the event LSTM performed on par with the Optical CNN with regards to accuracy ($\approx 92.5\%$). While it is less accurate than the event CNN and SNN, it is a smaller model than the CNN, with fewer model parameters as seen in Table 5.1. Furthermore, it performs less favorably than both the event CNN and SNN in terms of inference time.

The Fourier-based method, is by far the most time efficient, but struggled much more than the optical CNN in accuracy. This is likely due to the fewer features compared to the other models. This is expanded upon in Section 5.4. Due to the need to access the full size of training data for clustering, it also is the largest model, where the model size scales in proportion to the training data.

The confusion matrices for each of the models is shown in Figure 5.1, and the results are highlighted in Table 5.1. What follows is a more thorough examination of each technique and how they perform in flow regime classification. Additionally, further discussion of the results and how each model performs with regards to precision, recall, and F-1 score, are included in Appendix D.

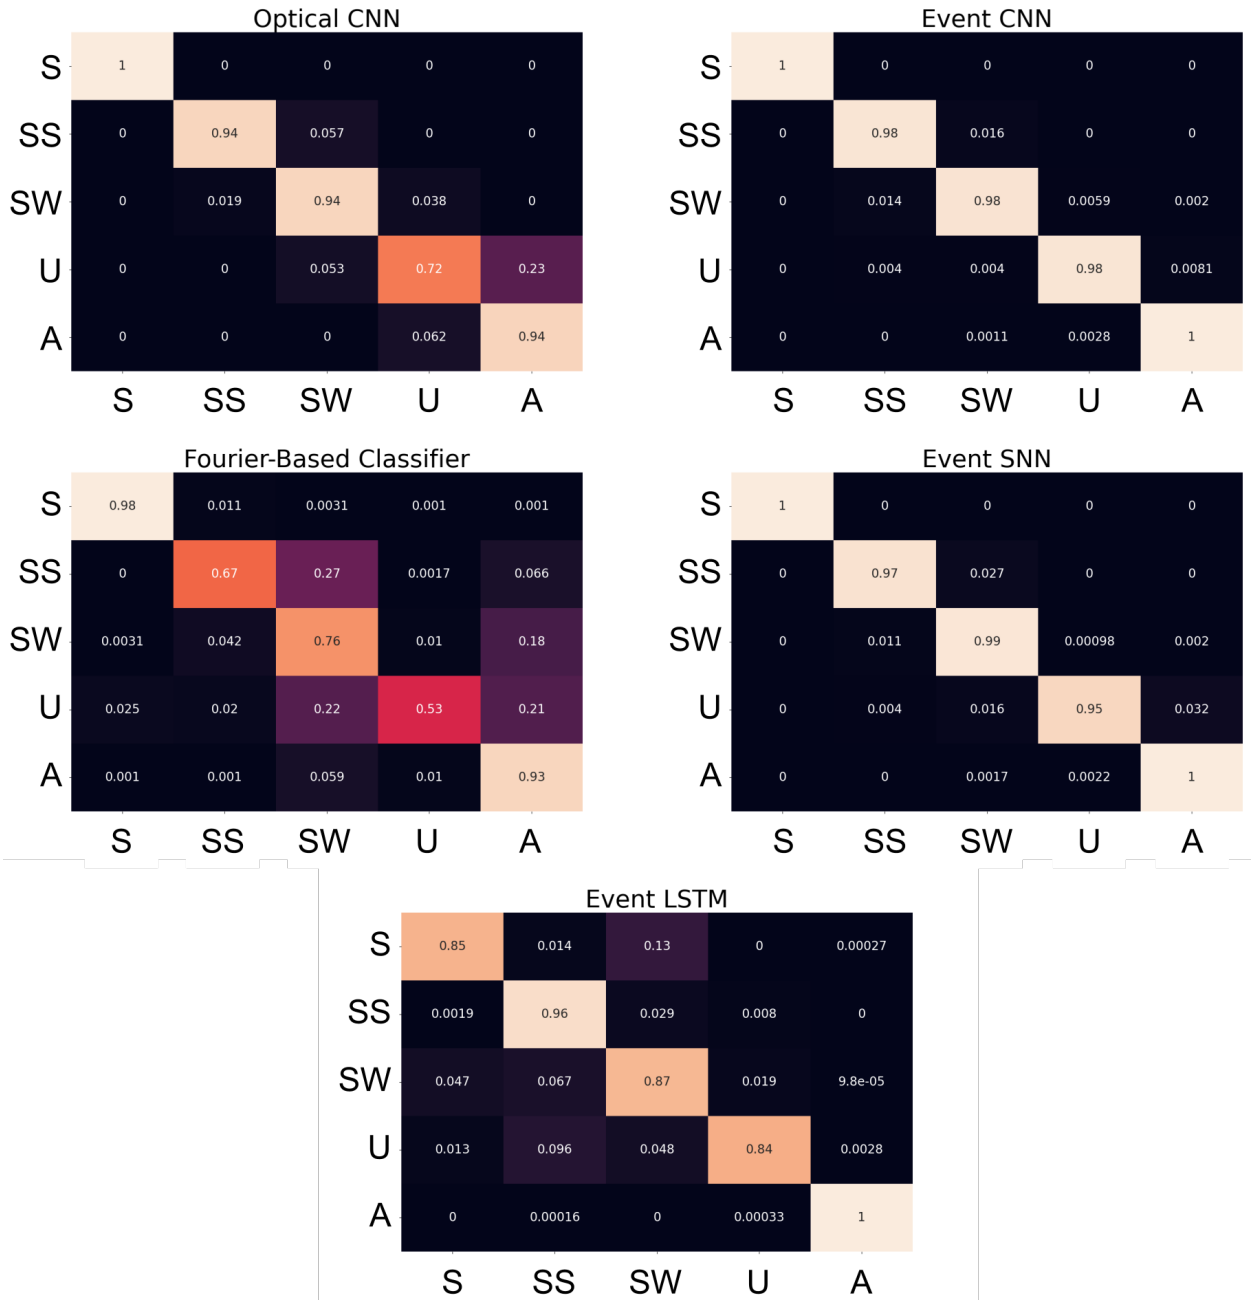

Figure 5.1: **Confusion Matrices across 5 classes.** Confusion Matrix for (a) Optical CNN, (b) Event CNN, (c) Fourier-based, (d) SNN, and (e) Event LSTM on test data. *A* represents Annular, *S* for Slug, *SS* for Stratified Smooth, *SW* for Stratified Wavy, and *U* for Unstable.

5.1 CNN

The Convolutional Neural Network is chosen because of both its ubiquity in classification tasks, but also their speed. In designing a model that would best tackle this difficult prob-

Model	Optical Dataset	Event Dataset	Event Model Latency	Event Model Size	Event Model Parameters
ConvNeXt (Scratch)	81.46%	72.26%	0.23ms	334.06MB	88.6M
ConvNeXt (Pretrained)	87.64%	88.73%	0.23ms	334.06MB	88.6M
AlexNet (Scratch)	87.01%	96.99%	0.17ms	217.53MB	61.1M
AlexNet (Pretrained)	84.79%	85.08%	0.17ms	217.53MB	61.1M
Custom	**92.24%**	**99.04%**	**0.15ms**	**1.90MB**	**0.5M**

Table 5.2: **Results comparing the custom Event and Optical CNN with ConvNeXt and AlexNet.** The pretrained models used weights trained on ImageNet data.

lem, multiple experiments are run to contrast different architectures. The results of these experiments are highlighted in Table 5.2. ConvNeXt, the current state of the art in image classification, is trained on both the optical and event data from scratch by loading the model from Pytorch with untrained weights. Another experiment is fine-tuning ConvNeXt on both the event and optical data. Training the model from scratch proved to be highly inaccurate, achieving an accuracy of around 81%. This is likely due to the relatively smaller size of the input data compared to the size of the model (Table 5.2). While ConvNeXt has shown to be very successful for highly detailed images, flow boiling images are not as varied or detailed. This likely contributes to the model not learning well. This is made apparent with the event dataset achieving an even lower accuracy of around 72%. The difference in accuracy is reflective of the event dataset being even more sparse than the already sparse optical dataset. However, fine-tuning the model with its existing ImageNet weights proves to increase the accuracy for both the optical and event data ($\approx$ 88% and 89% respectively). Though, now, the event data proves to be the more accurate model. This increase for both datasets is likely due to more of the neurons being active for classification and having substantive weights to them from ImageNet training. Another experiment, fine-tuning AlexNet in a similar way to ConvNeXt, performs slightly worse in accuracy at around 85%. However,

the model is smaller and has a better inference time as can be seen in Table 5.2. Finally, the shallow custom network performs far better than even ConvNeXt for this particular problem. A shallower network seems to better fit the low detail and sparse data when compared to the large networks. However, despite being a smaller model, the event CNN provides an almost 7% improvement in accuracy as compared to the optical CNN.

The Optical CNN is represented by the confusion matrix (a) in Figure 5.1, while the Event CNN is represented by the confusion matrix (b). While the accuracy may still be high for the optical CNN, the confusion matrix exposes both, that the stratified regimes remain to be a problem point for this network, but also the difficulties in categorizing unstable flow. This is a major concern, as these unstable regimes are both very uniquely interesting physically, but also a major area of concern for industrial systems. Flow instability can be a marker for a breakdown in a system's ability to properly perform heat management, as well as indicate pressure or temperature changes. Additionally, experiments into understanding the physics of unstable flow are hampered by the inability to extract unstable data from other regimes. The event CNN is able to better distinguish unstable flow than the optical CNN (Figure 5.1 (a) and (b)). This is likely due to the density of data in the event representation. The event data is binary, and instability results in a large clustering of data *exclusively* in the unstable region. The optical data meanwhile, captures both the unstable region as well as the rest of the channel (Figure 3.1). Unstable data is drawn from both Annular and Stratified regimes, and this confusion leads to the unstable data often being classified as being one of the regimes it came from. This is made obvious in the confusion matrix for optical data (Figure 5.1 (a)), highlighting that unstable regimes are classed as either Annular or Stratified Wavy around 30% of the time. While the optical CNN is only able to classify around 70% of unstable regime samples correctly, the event CNN performs at around 98%.

5.2 LSTM

The Event LSTM is represented by the confusion matrix (e) in Figure 5.1. It performs less accurate prediction in a longer time than the event CNN, but it does so with a smaller model. Interestingly, it performs best when tackling the *annular-stratified problem*, as seen in the high accuracy in both the annular and stratified wavy regimes. This is due to the tight and predictable clustering of events in these regimes. The slug and unstable regimes, while well predicted in other networks (see 5.1), are less accurately predicted in the LSTM. This network makes use of a series of x,y,p components of the event data, and thus a flowing series of events in these two regimes, can be mistaken for annular and stratified regimes. This is made apparent in the confusion matrix for the LSTM (Figure 5.1 (e)), wherein the slug and unstable regimes are confused for the annular and stratified regimes.

5.3 SNN

The Event Spiking Neural Network, as mentioned above, performs very similarly to the optical CNN. This makes sense, as the event SNN is a very similar architecture, although as mentioned in Section 4.4, the SNN consists of one less convolutional layer. One would expect a similar architecture to give similar results, though the SNN really shines in it's much smaller size for its accuracy, as seen in Table 5.1. The main draw to these networks are their speed, power consumption, and size. While the SNN did not perform as well as the event CNN in terms of it's latency, it did prove to be much faster than the optical CNN. Part of the latency issue for the SNN may be due to the fact that it is not run on neuromorphic hardware, to best utilize its sparsity. Additionally, SNN inference time is highly dependant on the number of spiking neurons per sample. Given the large number of accumulated events (50,000) per sample, a convolutional spiking neural network may produce a large number of

spiking events, thus increasing its latency. Regardless of hardware, shallow networks, such as the event SNN, have been shown to perform similarly to CNNs with regards to inference time. As networks get more complex, however, SNNs perform more favorably than CNNs with regards to latency [27]. The Convolutional SNN in particular has a large number of initial computations and thus would benefit from neuromorphic hardware [28].

Additionally, like the event CNN, it performed almost 7% better in terms of it's overall accuracy, but struggled similarly to the event CNN in the unstable regime. The SNN correctly classified 95% of the unstable regime as unstable, as opposed to the event CNNs 98% but the event SNN does a better job at distinguishing the extremely sparse stratified smooth regime as seen in the confusion matrix (d) of Figure 5.1.

5.4 Fourier Analytical Method

The Fourier-based approach is less accurate than the machine learning approaches. This analytical heuristic is able to achieve an almost 85% accuracy, however, further work, and better leveraging of signal processing techniques may help increase accuracy. It, expectedly, performed the best when accounting for inference time, but the poor accuracy may prove to render the latency gains less than useful. As seen in confusion matrix (c) of Figure 5.1, it does a good job of highlighting the slug regimes and annular regimes, but like the optical CNN, struggles with regards to the unstable regime. This is likely due to the physical properties of annular and slug regimes. In the annular regime, the thin layer of liquid around the vapor leads to lots of high frequency noise in both the x and y dimensions. This boundary region is subject to lots of small clusters of events moving through the channel. On the other hand, the slug regime consists of very little noise in the x and y direction due to the bubbles being clearly distinct from the surrounding liquid in the channel. The clear separation leads to much less noise as can be seen in the red cluster in Figure 4.3. This method takes these

frequency components into account, and thus these two regimes are made clear. On the other hand, the *annular-stratified problem* is readily apparent with this method. The three regimes are hard to distinguish between each other. The unstable regime likewise performs poorly when compared with the CNN and SNN models. This is possibly a result of unstable regimes containing signal characteristics similar to those of the regimes they originate in. For example, an annular flow that becomes unstable likely has a similar power spectrum or signal-to-noise ratio as a pure annular data point. This is reflected in many unstable flows being classified as either annular or stratified wavy, two regimes from which unstable flow originates. Additionally, this Fourier-based method processes the data and performs its heuristic on the entire frame. The CNN models, however, can take into account local dynamics of the data, and thus highlight the differences in spatial density of events or the lack of data in certain areas of the channel. Interestingly, despite having a higher recall score (0.67), the stratified smooth regime has a lower precision (0.54) than the unstable regime (0.53 and 0.81) respectively. This implies that this method is unlikely to classify a flow pattern as unstable, and more likely to consider it stratified smooth. The specifics of these metrics are further expanded on in Appendix D, but this suggests that this method struggles to create clusters for unstable flow patterns. While future work may elucidate better techniques, this model currently provides the lowest accuracy and largest size, two things that make it less useful for flow regime classification, than some of the previous models.

5.5 Seven Class Classification

Applying the same methods to the 7-class classification yields very similar results to the earlier 5-class classification. Like above, the CNN and SNN perform the best with respect to accuracy (each around 99%), however the SNN is half the size with half the parameters. The CNN, however, is slightly faster (Table 5.3). Additionally, the Fourier-based method

Model	Accuracy	Inference Time	Model Size	Model Parameters
Event CNN	99.051%	0.424ms	2.647MB	694,023
Event SNN	**99.090%**	0.875ms	**1.317MB**	**345,319**
Fourier	87.035%	**47.7**μs	7.338MB	N/A

Table 5.3: **Final results for 7 class flow regime classification.** Bold terms represent best performing metrics.

continues to perform the fastest.

In each case, the new regimes, Bubbly and Elongated Bubbly, are consistently classified with high accuracy and with very little confusion between them as seen in Figure 5.2. The misclassification among the stratified wavy, stratified smooth, and unstable regimes, however, continues to be a problem for the Fourier-based classifier. Regimes that occur early in the boiling process (Bubbly, Elongated Bubbly, and Slug) have very high recall, rarely being misclassified. However, in the middle regimes (Stratified Smooth, Stratified Wavy, and Unstable), the recall is lower, and the network struggles to accurately classify the data (Figure 5.2 (c)). Additional metrics such as precision, recall, and F-1 score are discussed in Appendix D.

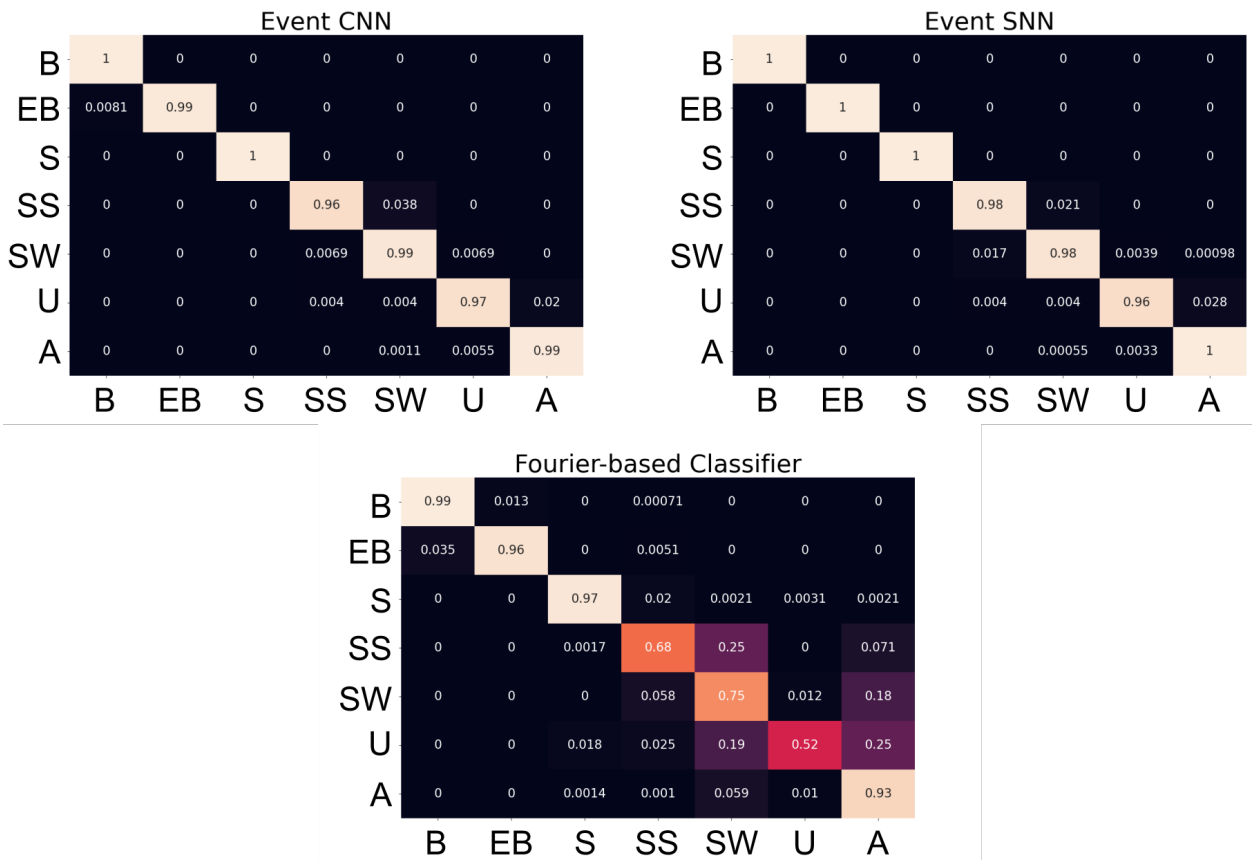

Figure 5.2: **Confusion Matrices across 7 classes.** Confusion Matrix for (a) Event CNN, (b) Event SNN, and (c) Fourier-based on test data. *A* represents Annular, *B* for Bubbly, *EB* for Elongated Bubbly, *S* for Slug, *SS* for Stratified Smooth, *SW* for Stratified Wavy, and *U* for Unstable.

5.6 Ablation

One confounding variable in comparing the size of the datasets. As noted in Table 3.1, there are significantly fewer optical frames than event frames. Some of the disparity between the accuracy of the optical and event CNNs may be due to the difference in the size of each dataset. In addressing this issue, an experiment is run on a reduced version of the event dataset to match the same exact class breakdown of the optical dataset. A random selection of each class in the event dataset was picked to match the number in the associated optical dataset. The event CNN was then run on this new, smaller dataset. The final accuracy for this new event CNN was **97.463%**. This is still a more accurate performance than the

optical CNN at **92.235%**. The resulting confusion matrices are shown in Figure 5.3. As shown in this figure, the event CNN continues to perform significantly better on discerning the unstable regime.

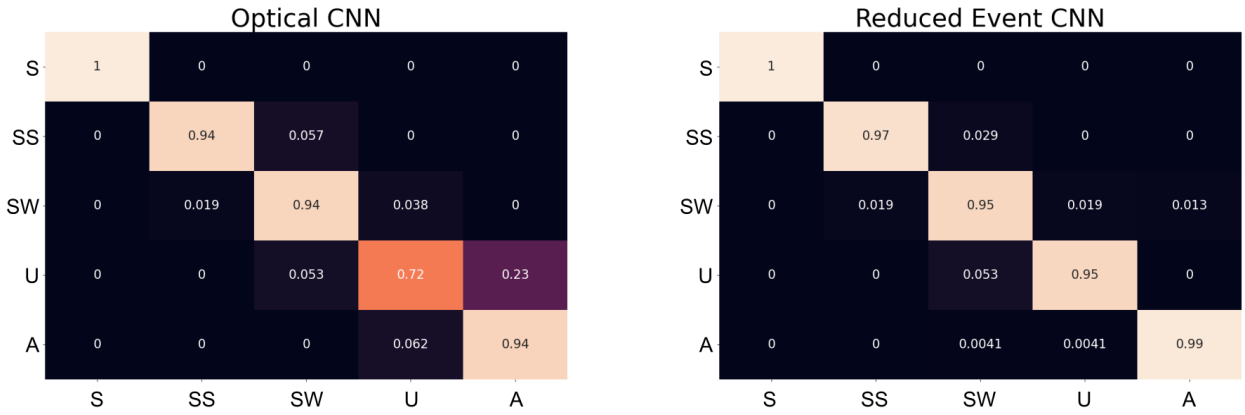

Figure 5.3: **Confusion Matrices for Reduced Event CNN as opposed to the Optical CNN**

Chapter 6

Conclusion

The significant accuracy and inference time benefits of event data in both CNNs and SNNs, as opposed to traditional optical CNNs, for flow regime classification present interesting implications for future experimentation and practical monitoring of industrial cooling systems. We present the case for event data models as a faster, more space efficient, and more accurate method for discerning flow patterns. Through an exploration of various machine learning techniques and heuristics, the goal of developing real-time regime classification models becomes more realistic. Further testing and exploration of the signal properties inherent in event data representations may additionally lead to a more unified definition on what constitutes a specific flow regime. Additionally, current classification problems dealing with unclear class distinctions may benefit from observing their classes with a new representation through event data. While the generalizeability of this approach may require future experimentation and work, the framework demonstrated here serves as a starting point for developing neuromorphic event models for other similar physics problems.

Links to download the code, as well as NeuroFlow optical and event datasets are available on the GitHub page, hosted at https://github.com/HPCForge/NeuroFlow.

Bibliography

[1] Mohamed H Mousa, Cheng-Min Yang, Kashif Nawaz, and Nenad Miljkovic. Review of heat transfer enhancement techniques in two-phase flows for highly efficient and sustainable cooling. *Renewable and Sustainable Energy Reviews*, 155:111896, 2022.

[2] Satish G Kandlikar. Heat transfer mechanisms during flow boiling in microchannels. *J. Heat Transfer*, 126(1):8–16, 2004.

[3] Biao Wang, Yanwei Hu, Yurong He, Nikolay Rodionov, and Jiaqi Zhu. Dynamic instabilities of flow boiling in micro-channels: A review. *Applied Thermal Engineering*, 214: 118773, 2022.

[4] Youngjoon Suh, Aparna Chandramowlishwaran, and Yoonjin Won. Artificial intelligence for liquid-vapor phase-change heat transfer. *arXiv preprint arXiv:2309.01025*, 2023.

[5] Eric Hervieu. Identification of gas-liquid flow regimes from a space-frequency representation by use of an impedance probe and a neural network. In *Fluids Engineering Division Summer Meeting*, volume 36169, pages 685–691, 2002.

[6] Marcelo Fernando Selli and Paulo Seleghim Jr. On-line identification of horizontal two-phase flow regimes through gabor transform and neural network processing. In *International Pipeline Conference*, volume 42630, pages 813–820, 2006.

[7] Feng Nie, Haocheng Wang, Qinglu Song, Yanxing Zhao, Jun Shen, and Maoqiong Gong. Image identification for two-phase flow patterns based on cnn algorithms. *International Journal of Multiphase Flow*, 152:104067, 2022.

[8] MK Seal, SMA Noori Rahim Abadi, M Mehrabi, and JP Meyer. Machine learning classification of in-tube condensation flow patterns using visualization. *International Journal of Multiphase Flow*, 143:103755, 2021.

[9] Dale Lu, Youngjoon Suh, and Yoonjin Won. Rapid identification of boiling crisis with event-based visual streaming analysis. *Applied Thermal Engineering*, 239:122004, 2024.

[10] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180, 2020.

[11] Cedric Scheerlinck, Henri Rebecq, Daniel Gehrig, Nick Barnes, Robert Mahony, and Davide Scaramuzza. Fast image reconstruction with an event camera. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 156–163, 2020.

[12] Yin Bi, Aaron Chadha, Alhabib Abbas, Eirina Bourtsoulatze, and Yiannis Andreopoulos. Graph-based object classification for neuromorphic vision sensing. *CoRR*, abs/1908.06648, 2019. URL `http://arxiv.org/abs/1908.06648`.

[13] X. Huo, L. Chen, Y.S. Tian, and T.G. Karayiannis. Flow boiling and flow regimes in small diameter tubes. *Applied Thermal Engineering*, 24(8):1225–1239, 2004. ISSN 1359-4311. doi: https://doi.org/10.1016/j.applthermaleng.2003.11.027. URL `https://www.sciencedirect.com/science/article/pii/S1359431104000250`. The 8th UK National Conference on Heat Transfer.

[14] Leonor Hernández, J. Enrique Julia, Basar Ozar, Takashi Hibiki, and Mamoru Ishii. Flow Regime Identification in Boiling Two-Phase Flow in a Vertical Annulus. *Journal of Fluids Engineering*, 133(9):091304, 09 2011. ISSN 0098-2202. doi: 10.1115/1.4004838. URL `https://doi.org/10.1115/1.4004838`.

[15] Zhiee Jhia Ooi, Longxiang Zhu, Joseph L. Bottini, and Caleb S. Brooks. Identification of flow regimes in boiling flows in a vertical annulus channel with machine learning techniques. *International Journal of Heat and Mass Transfer*, 185:122439, 2022. ISSN 0017-9310. doi: https://doi.org/10.1016/j.ijheatmasstransfer.2021.122439. URL `https://www.sciencedirect.com/science/article/pii/S0017931021015374`.

[16] Gustavo M. Hobold and Alexandre K. da Silva. Machine learning classification of boiling regimes with low speed, direct and indirect visualization. *International Journal of Heat and Mass Transfer*, 125:1296–1309, 2018. ISSN 0017-9310. doi: https://doi.org/10.1016/j.ijheatmasstransfer.2018.04.156. URL `https://www.sciencedirect.com/science/article/pii/S0017931017346100`.

[17] Alberto Sabater, Luis Montesano, and Ana C. Murillo. Event transformer$^+$. a multi-purpose solution for efficient event data processing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(12):16013–16020, 2023. doi: 10.1109/TPAMI.2023.3311336.

[18] Udayanga K. N. G. W. Gamage, Luca Zanatta, Matteo Fumagalli, Cesar Cadena, and Silvia Tolu. Event-based classification of defects in civil infrastructures with artificial and spiking neural networks. In Ignacio Rojas, Gonzalo Joya, and Andreu Catala, editors, *Advances in Computational Intelligence*, pages 629–640, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43078-7.

[19] Chang Liu, Xiaojuan Qi, Edmund Y. Lam, and Ngai Wong. Fast classification and action recognition with event-based imaging. *IEEE Access*, 10:55638–55649, 2022. doi: 10.1109/ACCESS.2022.3177744.

[20] Bochen Xie, Yongjian Deng, Zhanpeng Shao, Hai Liu, and Youfu Li. Vmv-gcn: Volumetric multi-view based graph cnn for event stream classification. *IEEE Robotics and Automation Letters*, 7(2):1976–1983, 2022. doi: 10.1109/LRA.2022.3140819.

[21] Teresa Serrano-Gotarredona and Bernabé Linares-Barranco. Poker-dvs and mnist-dvs. their history, how they were made, and other details. *Frontiers in Neuroscience*, 9, 2015. ISSN 1662-453X. doi: 10.3389/fnins.2015.00481. URL `https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2015.00481`.

[22] Hongmin Li, Hanchao Liu, Xiangyang Ji, Guoqi Li, and Luping Shi. Cifar10-dvs: An event-stream dataset for object classification. *Frontiers in Neuroscience*, 11, 2017. ISSN 1662-453X. doi: 10.3389/fnins.2017.00309. URL `https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2017.00309`.

[23] Garrick Orchard, Ajinkya Jayawant, Gregory K. Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in Neuroscience*, 9, 2015. ISSN 1662-453X. doi: 10.3389/fnins.2015.00437. URL `https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2015.00437`.

[24] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s, 2022.

[25] Jason K Eshraghian, Max Ward, Emre Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Bennamoun, Doo Seok Jeong, and Wei D Lu. Training spiking neural networks using lessons from deep learning. *Proceedings of the IEEE*, 111(9): 1016–1054, 2023.

[26] Ali Samadzadeh, Fatemeh Sadat Tabatabaei Far, Ali Javadi, Ahmad Nickabadi, and Morteza Haghir Chehreghani. Convolutional spiking neural networks for spatio-temporal feature extraction, 2021.

[27] Patrick Plagwitz, Frank Hannig, Jürgen Teich, and Oliver Keszocze. To spike or not to spike? a quantitative comparison of snn and cnn fpga implementations, 2023.

[28] Pao-Sheng Vincent Sun, Alexander Titterton, Anjlee Gopiani, Tim Santos, Arindam Basu, Wei D. Lu, and Jason K. Eshraghian. Intelligence processing units accelerate neuromorphic learning, 2022.

[29] Yuhuang Hu, Shih-Chii Liu, and Tobi Delbrück. V2E: from video frames to realistic DVS event camera streams. *CoRR*, abs/2006.07722, 2020. URL `https://arxiv.org/abs/2006.07722`.

Appendix A

Neuromorphic Event Data

The development of a circulating loop system for flow boiling has been accomplished, incorporating a milled polycarbonate rectangular internal channel and a copper-based resistive heater, as illustrated in Figure A.1. The rectangular flow channel is characterized by a cross-section of 2.5 by 5 mm, a hydraulic diameter of 3.33 mm, and a total length of 380 mm, allowing for single-sided heating on the bottom wall. The transparent channel wall, with a thickness of 13.75 mm, facilitates clear visualization of flow bubbles. To enable a wider range of boiling operations, dielectric working fluid (PP1) with lower saturation temperatures has been selected. The 100 mm heated section of the channel is powered by six thick film resistors, each possessing a resistance value of 200 ohms. These resistors are affixed to the plate and wired in parallel, resulting in a cumulative resistance of 33.33 ohms. A controllable power source facilitates precise modulation of the heat flux applied to the system.

Ensuring optimal operational conditions, a liquid reservoir equipped with a circulation heater preheats the working fluid to near-saturation temperatures before entering the main channel. Subsequently, the preheated fluid is pumped to a flow meter, allowing for meticulous control of the flow rate before entering the experimental channel. The data acquisition apparatus

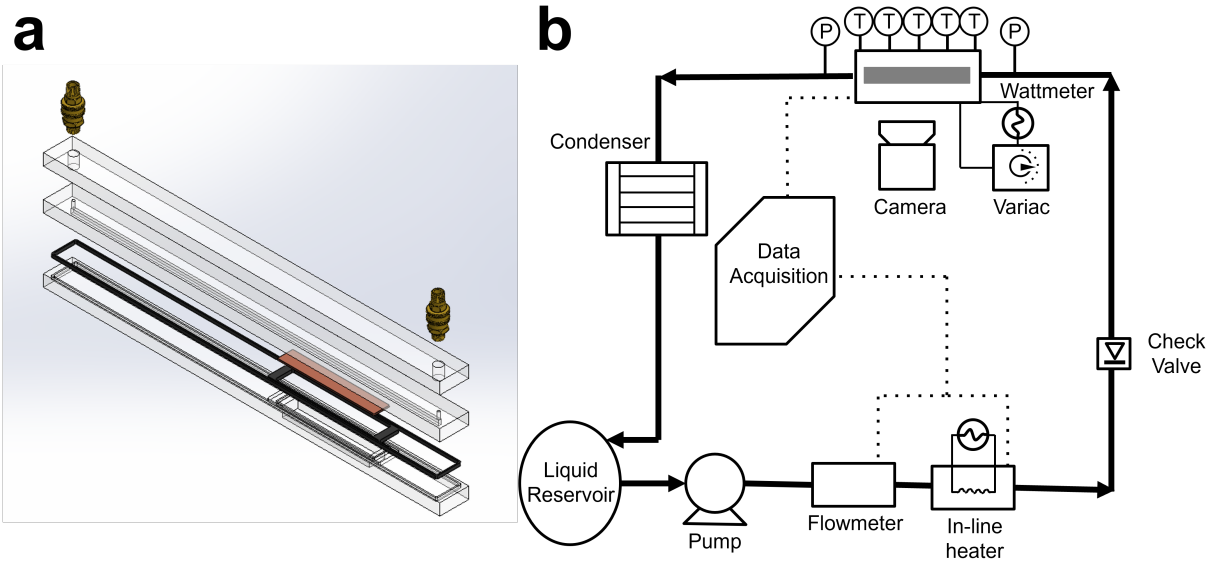

Figure A.1: **Visualization for flow boiling.** (a) The flow boiling module consists of a closed-loop system. (b) The flow diagram of the flow boiling experimental set-up.

collects pressure and temperature data at the channel's inlet and outlet. Upon exiting the flow channel, the working fluid enters a condenser, thereby completing the flow loop, as depicted in Figure A.1.

To capture diverse visual aspects of the flow boiling scenario, we leverage an event-based neuromorphic camera. Traditional high-speed cameras store densely packed data with comprehensive information, offering precise representations of visual scenes at high resolutions. However, a substantial portion of this information may be redundant, leading high-speed imaging methods to have constrained run times. In contrast, event cameras utilize innovative sensors that report asynchronous, per-pixel brightness changes termed as "events" with unparalleled low latency.

By exclusively recording brightness changes, event cameras efficiently eliminate redundant information, such as background details, between time intervals without the need for pre- or post-processing steps. When a pixel detects a change in light intensity, it autonomously triggers and records the event as depicted in Figure A.2. Furthermore, modern event cameras

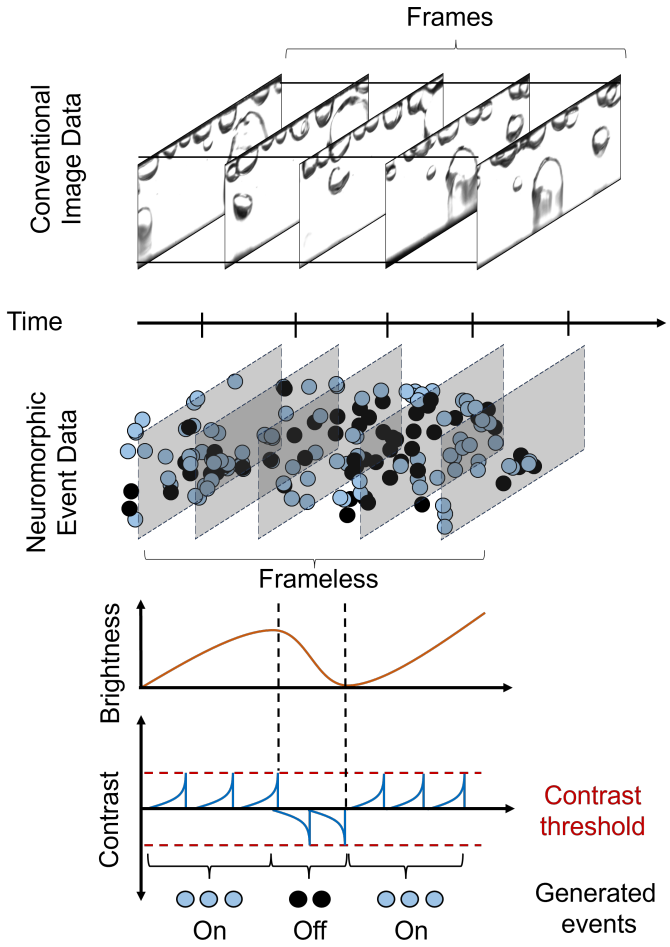

Figure A.2: **Comparison between conventional image data and neuromorphic event data.** Event data prioritizes scene changes, disregarding static backgrounds, offering a more efficient and responsive method for capturing visual information than conventional cameras.

exhibit microsecond temporal resolution, minimal under/overexposure, and motion blur, enhancing their capability to rapidly capture dynamic motions. Given that high-speed camera runtimes are often restricted by high power consumption, the significant increase in power efficiency with neuromorphic cameras suggests prolonged continuous streaming of data is feasible.

Appendix B

V2E Conversion

V2E, or Video to Events, is an advanced Python software tool meticulously crafted for the conversion of conventional video frames into authentic Dynamic Vision Sensor (DVS) event streams. It deftly generates DVS event camera data, capturing nuances in scene brightness, from both standard and synthetic frame-based videos. This process relies on a precise DVS pixel model, accounting for the inherent non-idealities of the camera. Additionally, V2E has the capability to employ Super-SloMo synthetic slow motion for enhancing the quality of videos from standard frame cameras through up-sampling. The tool finds its primary application in the realms of computer vision and neuromorphic engineering research, where the high temporal resolution and low power consumption of DVS data hold significant value. Notably, V2E plays a pivotal role in creating transfer-learning training and evaluation datasets specifically tailored for event cameras, derived from conventional frame-based datasets [29].

This capability enables us to extract meaningful insights, promoting a proactive understanding and response to dynamic visual content. The integration of V2E enhances the utilization of video data, aligning with the evolving demands of modern technological landscapes for a more refined and sophisticated approach. We integrate this tool into the transfer-learning

process, using it to prepare training and evaluation datasets for event cameras by adapting them from conventional frame-based datasets. A total of 15 videos, recorded using high-speed cameras, have been transformed into event streams for in-depth analysis.

Appendix C

Experimental Data Tables

Prophesee Experiments			
Voltage	Total Events	Event Frames	Main Regime
25	1.63×10^8	3,266	Bubbly
30	6.83×10^7	1,367	Bubbly
35	7.39×10^7	1,479	Bubbly
40	6.03×10^7	1,206	Bubbly
45	7.20×10^7	1,441	Bubbly
55	6.17×10^7	1,235	Elongated Bubble

Table C.1: **Six experiments from which the Prophesee camera data is created.** In each of these experiments, only one regime was noted.

Emulated Data				
Mass Flux	Heat Flux	Total Events	Event Frames	Regimes
140	9.4	3.70×10^7	740	**Strat. Smooth**
	18	8.63×10^7	1,726	**Strat. Wavy** Strat. Smooth
	30.3	9.30×10^7	1,860	**Annular** Unstable
	43	8.79×10^7	1,757	**Annular**
200	6.8	1.15×10^8	2,308	**Slug**
	13	1.28×10^8	2,557	**Strat. Wavy** Strat. Smooth Unstable
	22	1.98×10^8	3,950	**Strat. Wavy** Strat. Smooth Unstable
	36.3	1.74×10^8	3,474	**Annular** Unstable
	53	1.51×10^8	3,029	**Annular** Unstable
	70	1.30×10^8	2,606	**Annular** Unstable
255	8	1.89×10^8	3,773	**Slug**
	15.4	1.38×10^8	2,767	**Strat. Wavy** Strat. Smooth Unstable
	28	1.80×10^8	3,609	**Strat. Wavy** Strat. Smooth Unstable
	48	1.70×10^8	3,408	**Annular** Unstable
	69.5	1.47×10^8	2,944	**Annular** Unstable

Table C.2: **Fifteen experiments from which the emulated event data is created.** In each of these experiments, different regimes are observed, the dominant regime of which, is in **bold**.

Appendix D

Precision, Recall, and F-1 Metrics

While traditional accuracy is a good measure for understanding a model's performance, it is incomplete. Accuracy is simply a measure of how many total samples are predicted correctly out of a total number of test samples. However, this does nothing to explore the ways in which a model may provide false positives or false negatives. Exploring these topics further leads to the metrics of precision, recall, and F-1 score.

D.1 Precision

Precision is a measure of how many correctly classified samples there are, relative to the total number of samples classified as such. Essentially, precision measures the ability of a model to accurately identify an instance of a class. A low precision score indicates that a model has a hard time distinguishing between various classes, and predicts a particular class regardless. This is an especially important problem when the datasets are rather unbalanced as they are in this problem. The formula for precision is given below.

$$\textbf{Precision}_{class\ A} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}}$$

D.2 Recall

Recall is a measure of how many samples of a particular class are correctly identified as being in that class. This is identical to the diagonal of the confusion matrix. Essentially, Recall provides the models ability to distinguish all samples of a particular class from other classes. A low precision score may indicate that a model is struggling to identify a particular class in general. The formula for recall is given below.

$$\textbf{Recall}_{class\ A} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}}$$

D.3 F-1 Score

The F-1 score is a compromise between the two metrics of precision and recall. It takes both the false positives and false negatives into account, and is a good measure for the accuracy of an unbalanced dataset like the one worked with here. The formula for F-1 Score is given below.

$$\textbf{F1}_{class\ A} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

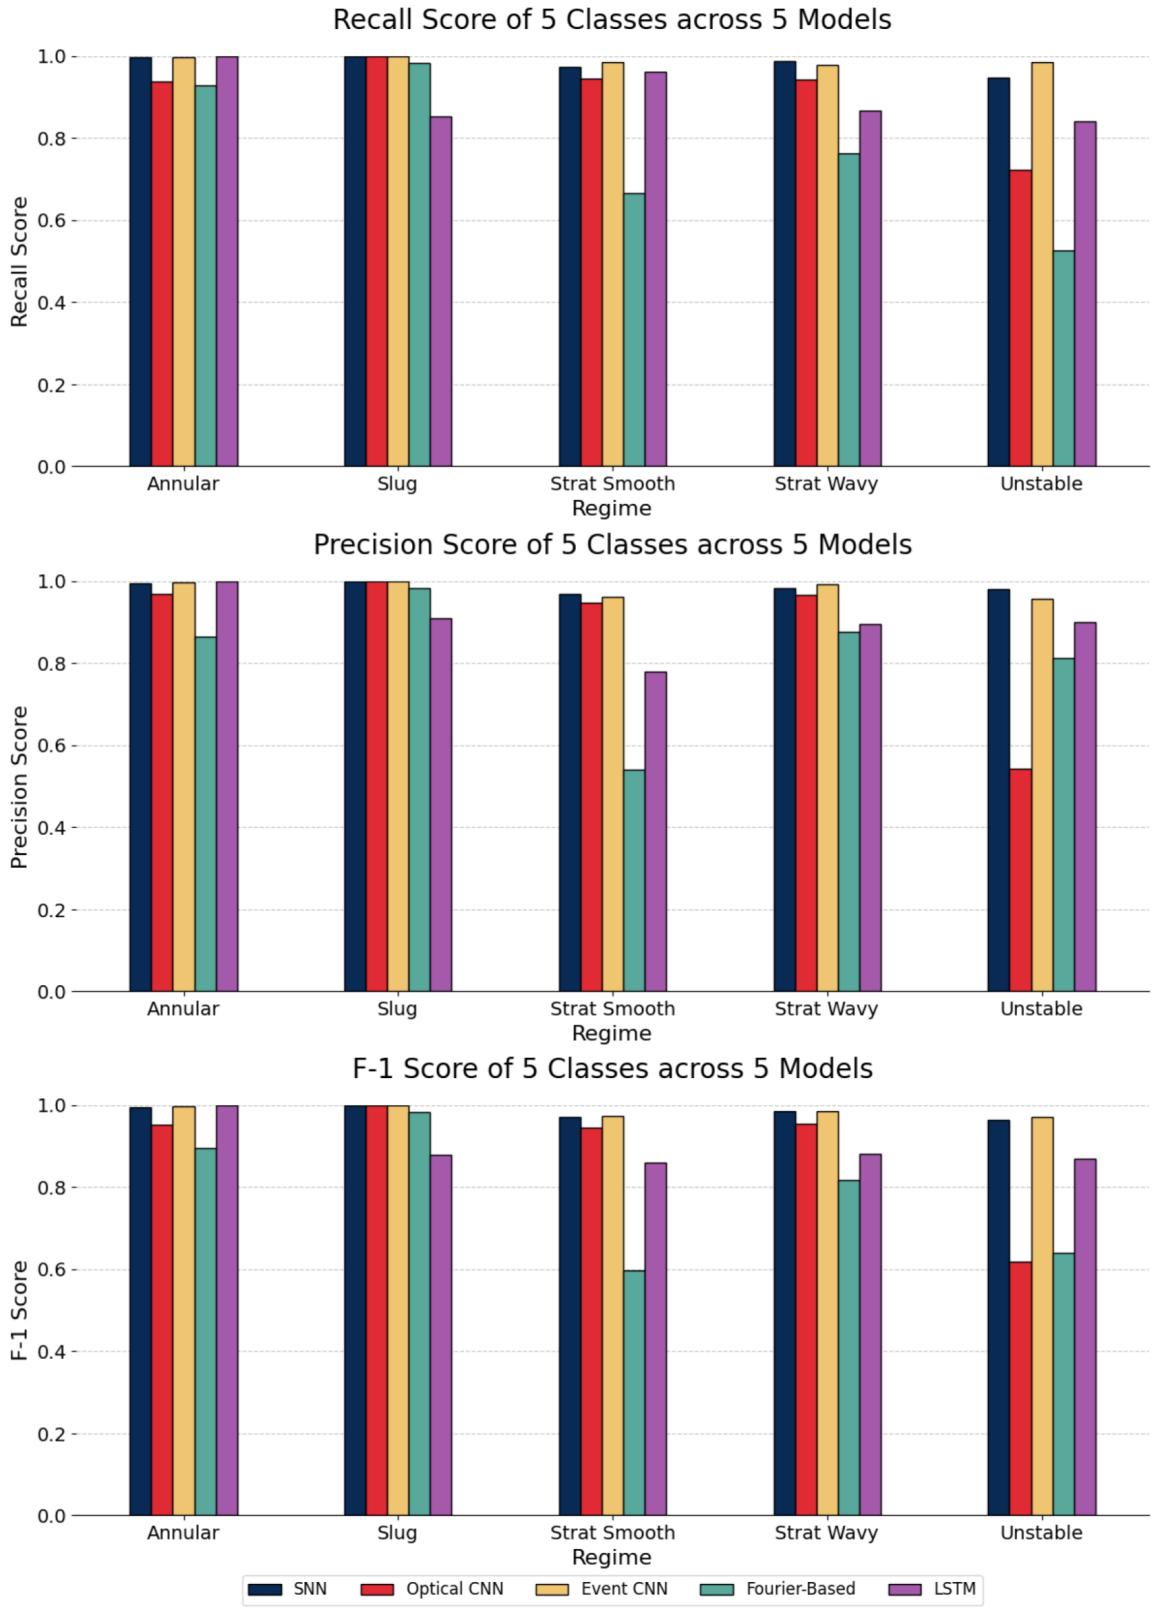

Figure D.1: **Precision, Recall, and F-1 Scores for the five compared methods, SNN, Optical CNN, Event CNN, Event LSTM, and Fourier-based Method.**

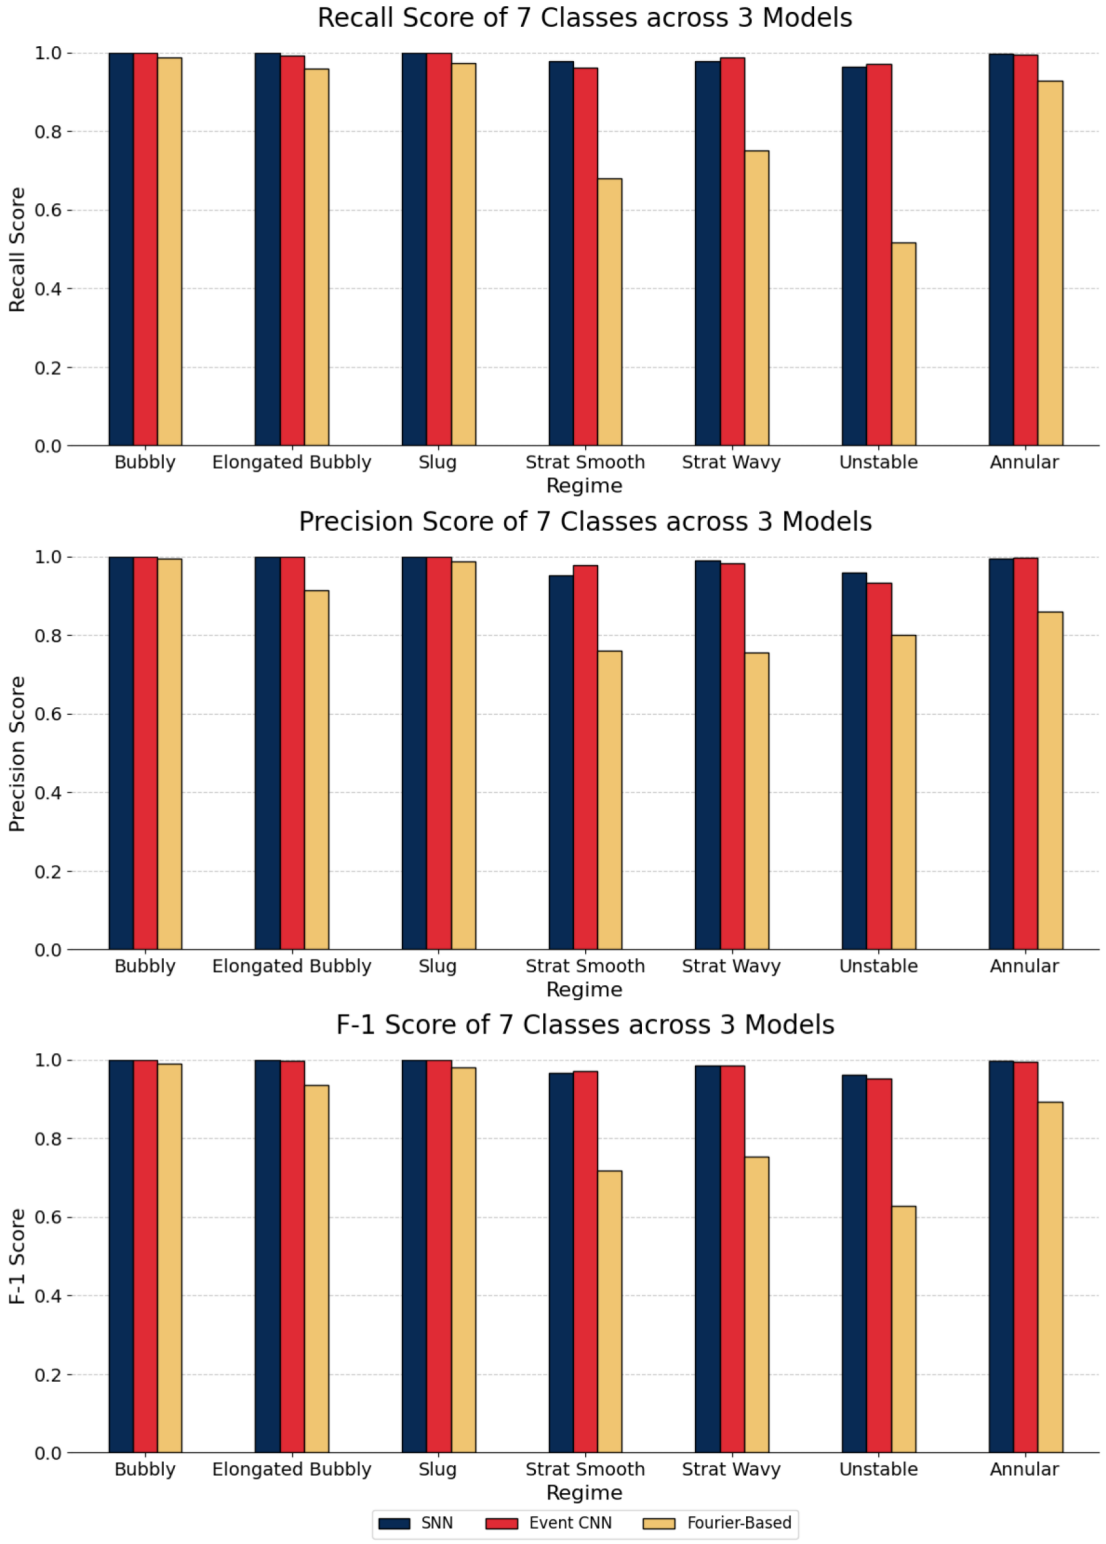

Figure D.2: **Precision, Recall, and F-1 Scores for the seven regimes on SNN, Event CNN, and Fourier-based Method.**