Applied Thermal Engineering 239 (2024) 122004

Contents lists available at ScienceDirect

Applied Thermal Engineering

journal homepage: www.elsevier.com/locate/apthermeng

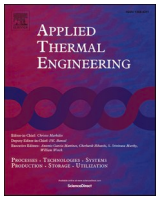

Research Paper

Rapid identification of boiling crisis with event-based visual streaming analysis

Dale Lu [a], Youngjoon Suh [b], Yoonjin Won [b,c,*]

[a] Department of Computer Science, University of California, Irvine, 3019 Donald Bren Hall, Irvine, CA 92697-3425, USA
[b] Department of Mechanical and Aerospace Engineering, University of California, Irvine, 4200 Engineering Gateway, Irvine, CA 92697-2700, USA
[c] Department of Electrical Engineering and Computer Science, University of California, Irvine, 2200 Engineering Hall, Irvine, CA 92697-2625, USA

ARTICLE INFO

Keywords:
Neuromorphic
Event
LSTM
Critical heat flux
Boiling

ABSTRACT

The formation of a stable vapor layer during boiling processes triggers the boiling crisis that can damage the entire system. While there have been efforts to build a predictive model for critical heat flux on live images, unfortunately, images from conventional frame-cameras show limitations for performing real-time digital inference as they contain significant unnecessary information. A new type of imaging technique, neuromorphic event imaging, captures individual pixel-wise brightness changes, thereby recording motion-based phenomena in a sparse and efficient manner. In this study, we develop neural network models capable of predicting the onset of boiling crisis by utilizing neuromorphic-simulated event data. We reveal that event-based models demonstrate their proficiency and time efficiency in learning patterns of the boiling process from relatively small datasets. The event-based models also prove the generalizability to unseen surfaces by showing the potential for transfer learning. The results suggest that an extension of this work may prove effective in developing forecasting models that detect the incipient boiling crisis before it damages boiling surfaces. These findings highlight that event data can be effectively applied to real-time, motion-dependent contexts.

1. Introduction

As a highly effective method of heat transfer, boiling is an appealing choice of temperature control in high-power systems [1] However, the efficiency of boiling is significantly reduced when nucleate boiling transitions into film boiling, also known as the critical heat flux (CHF). In addition to the deterioration of heat transfer, the onset of CHF induces a vapor blanket that insulates the heated surface from the working fluid, thereby causing extreme temperature surges which threatens the integrity of safety–critical systems relying on consistent temperatures [2].

Despite an extensive history of prior efforts to prevent and mitigate the damages caused by CHF, predicting the chaotic boiling behaviors remains a daunting challenge [3]. Historically, the prediction and forecasting of CHF have been a primary subject of interest in the thermal science community, where most of the approaches were based on mechanistic modeling [4,5]. However, due to the highly intertwined parametric factors that decide the boiling outcome, most modeling approaches are still limited to narrow and specific operating conditions [6,7]. As a consequence, the ultimate goal of CHF forecasting is still being developed from the relatively simpler task of CHF detection. Considering the important implications of this task, there are still surprisingly few methods that can accurately detect the occurrence of CHF in real-time [1,8–10]. One method to detect CHF is to collect temperature data from thermal probes or sensors [1,11]. Since the onset of CHF is quickly followed by a temperature surge, thermal measurements are by far the most accessible and straightforward way to identify CHF. However, such probing methods are poor at portraying the central nucleation activities that mechanistically lead to the CHF, and are therefore limited for the development of CHF forecasters [12]. Another alternative is using live images and domain knowledge to detect CHF [13]. For example, existing theory states that CHF is a state of heightened bubble–bubble interactions and growth, that leads to the systematic failure of the surface to relieve itself from the generated vapor [2]. Therefore, this investigation suggests that such visual cues in images can be used to predict CHF through machine learning. However, training a model to distinguish the visual patterns that identify CHF is, in practice, an immense challenge [6,14].

While there have been previous attempts using vision-based frameworks to automate CHF detection, [1,6,10,14] the use of conventional

* Corresponding author.
E-mail address: won@uci.edu (Y. Won).

https://doi.org/10.1016/j.applthermaleng.2023.122004
Received 26 May 2023; Received in revised form 11 November 2023; Accepted 12 November 2023
Available online 23 November 2023
1359-4311/© 2023 Elsevier Ltd. All rights reserved.

frame cameras for data collection has imposed significant limitations on the efficiency and real-time applicability of such studies [15,16]. Traditional frame-based cameras store data for every pixel, often in multiple color channels, with many pixels encoding redundant information (i.e., background data or noise) [17]. Accordingly, the captured image data are densely packed with information, but a significant proportion of this information is irrelevant for performing digital inference tasks. Coincidentally, high-speed (~1000–100,000 fps) cameras are customarily used to visualize rapid boiling processes, which typically occur in a few tens of microsecond cycles, with adequate resolutions, which further aggravates the data intensity [1,3]. Since the high data intensity of images typically correlates to larger computational expense (i.e., prolonged prediction times), prior ML researchers have proposed using preprocessing methods such as principal component analysis (PCA) or downscaling to reduce their dimensionality [6,10]. Yet, these approaches often underreport the '*actual time*' it takes to gather visual data and preprocess them, which is crucial for real-time predictions. In reality, high-speed cameras are inappropriate for real-time streaming applications because of their short (typically limited to a few seconds) runtimes. Thus, even if data storage and preprocessing steps were to be adequately shortened, they would still ultimately be workarounds for an inherently constraining hardware choice.

Both the data inefficiency and power consumption of traditional imaging techniques can be circumvented with an emerging alternative imaging method. Neuromorphic event cameras are shutterless paradigm-shifting sensors that report asynchronous, per-pixel brightness changes called "events" with unparalleled low latency [18] (Fig. 1). Conventional frame cameras operate by capturing individual frames at discrete times, determined by the camera's framerate. However, for images that capture boiling phenomena, only a portion of the encoded data captures relevant information for CHF detection, as seen in Fig. 1a. By only recording brightness changes, event cameras effectively ignore repeated information between time intervals, thereby removing redundancies like background information from the generated data without relying on preprocessing steps. Event cameras are therefore appropriate for recording motion without extraneous data, which encodes information like nucleation behavior within the context of CHF detection. This enables uniquely compact representations of visual information, holding the potential to facilitate cost- and energy-effective predictions when interpreting transient phenomena. Additionally, modern event cameras have microsecond temporal resolution, minimal under/overexposure, and motion blur, which lends them more potential for rapidly capturing dynamic motions [18]. In contrast to the high-speed cameras which have been observed to consume approximately 65 W, [3] some neuromorphic event cameras consume less than 10mW while using 10x less memory [11,19–21]. Since high-speed camera runtimes are largely limited by their high power consumption or memory usage, the neuromorphic cameras' notable advantages will allow for extended, continuous streaming of data, making them attractive options for real-world scenarios. Therefore, integrating data generated by event cameras into neural network architectures might lead to practical improvements in prediction speed while avoiding the runtime limitations imposed by frame-based cameras.

In this work, we present a deep learning approach for CHF that leverages the inherent efficiencies of neuromorphic event data to overcome the limitations of conventional image-based classification. Event data is simulated from pre-collected experimental image data to ensure the same phenomena is captured in both event and image formatting,

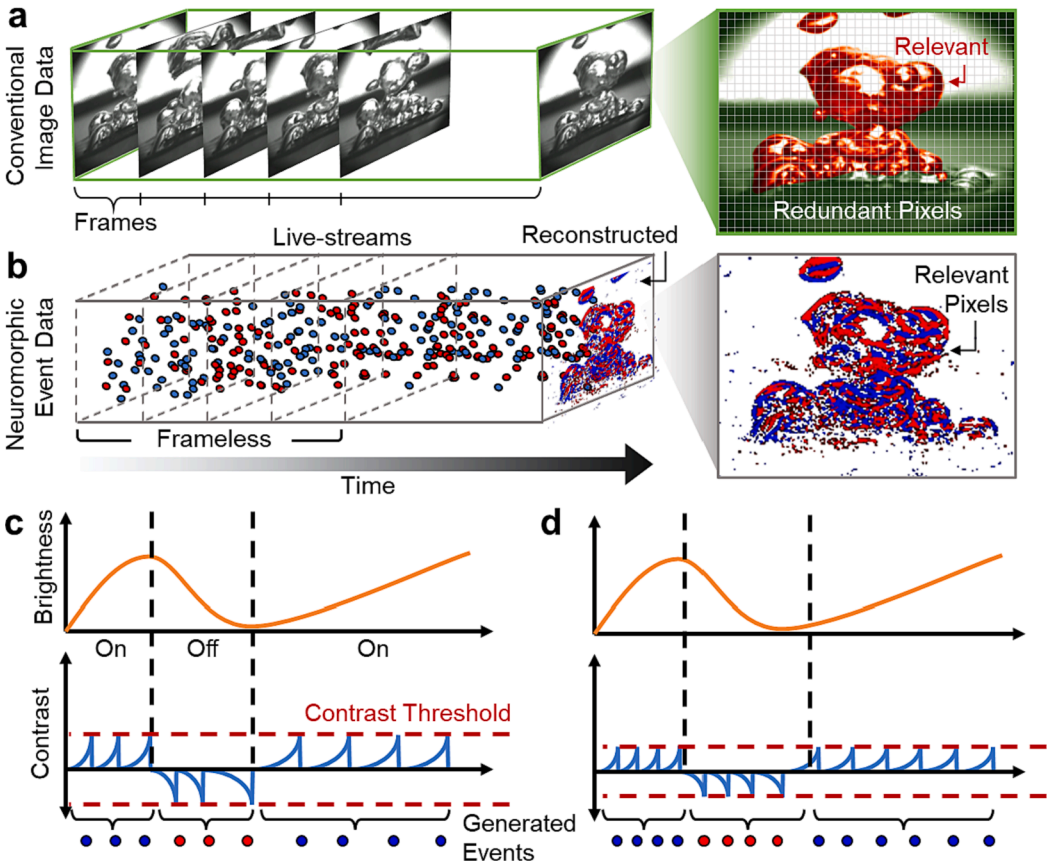

Fig. 1. (a) Conventional imaging captures data at discrete time intervals as well as redundant information. (b) Neuromorphic event cameras capture a stream of individual pixel-wise brightness changes, which effectively captures motion. (c) Neuromorphic cameras function by setting thresholds on each pixel's brightness change can detect excessive contrast as an event. (d) Changing parameters like contrast threshold will affect what constitutes an event as well as how many events are present.

and these data are distributed to various models for training purposes. By evaluating a wide array of traditional CNN-based classifiers and specialized event-based models on both their predictive capabilities and prediction times, we conduct a rigorous investigation into the advantages provided by event data. Our findings reveal that an event-based classification model can achieve over three times faster prediction speed than that of conventional image-based models, while demonstrating predictive capabilities on par with them.

2. Methodology

2.1. Experimental setup

All data is obtained from a pool boiling heat transfer setup composed of a chamber, data acquisition system, copper block, and heaters. Four cylindrical cartridge heaters (Omega, CIR-20191) positioned at the bottom of the copper block heat the samples. The heat flux is controlled by slowly increasing power using a voltage transformer (Variac AC variable voltage converter). Surface temperatures are recorded on a data acquisition system (LabJack U6) connected to four K-type thermocouples embedded in the copper block. Before the boiling test, the sample is soldered onto the copper block. A guard heater connected to a proportional integral derivative (PID) controller is used to maintain the working fluid at saturation temperature. Further description of the experimental setup can be found in previous work [1].

Side-view visualizations of bubble dynamics are captured through the transparent polycarbonate boiling chamber wall using a high-speed camera (FASTCAM Mini AX50) at 2,000 fps with an LED lighting source. The data is saved by skipping every 4 frames to reduce the overall image inventory size. To mitigate errors caused by different focal lengths, the camera and light source are fixated throughout all experiments. A light diffuser is installed between the lighting source and boiling chamber for even lighting.

Videos taken by the high-speed camera during pool boiling experiments are organized by their corresponding heat flux values. One video consisting of 2,184 frames is generated from CHF experiments. To have a dataset equally weighted between CHF and non-CHF videos, 21 smaller videos of 104 frames are compiled from discrete heat flux values ranging from 25 W/cm^2 to 390 W/cm^2. Event data simulation is performed by passing videos to the software tool *v2e*. *v2e* begins by converting frames of color video to luma frames, which only convey brightness [22]. Frame interpolation is performed to increase spatial and temporal resolutions before brightness changes are measured between frames to determine the time and position of events that would be generated by physical event cameras [22]. These events are stored in .aedat files, which efficiently store several millions of events per heat flux step.

Given the variety of models tested in this work, multiple preprocessing approaches are necessary to generate appropriate data formats. For a conventional image-based CNN, the side view visualizations are resized to a standardized input resolution. For CNNs predicting on images reconstructed from event data, events are accumulated into time intervals and mapped to a 2D surface, with polarity encoded into the pixel color. This results in an image reconstruction of the event data which only includes information on bubble motion. Similar images are used as inputs to the ConvLSTM model, but sequential images are grouped into short videos of consistent length. These short videos are compatible with the time-series data format that the ConvLSTM layer expects. For the eventLSTM, the image medium is ignored in favor of splitting the event dataset into intervals that emulate the timesteps necessary for LSTM inputs.

2.2. Model architectures

To systematically investigate how input data conveying varying contextual information can influence prediction accuracy, speed, and model generalizability, we train five neural network architectures on two broad categories of data: images and event strings. To ensure that both categories of data encode identical imaging conditions (i.e., viewing angle, duration, lighting), we convert acquired high-speed pool boiling images into simulated event data using an open-source conversion software *v2e* [22]. The conversion process emulates the working mechanisms of neuromorphic cameras by detecting brightness changes within the sequential images, thereby drastically reducing redundant pixel information. The *v2e* script additionally incorporates frame interpolation, [23] which further mimics the temporal resolutions of actual neuromorphic cameras. Therefore, the event i is structured to show four values (x_i, y_i, t_i, p_i): x and y values, the timestamp t, and the polarity of the event p.

CNN-Control. To compare our work with conventional image-based approaches, high-speed images are directly used as an input for a simple CNN architecture (i.e., CNN-Control) designed using similar components to prior work. To emulate prior work, these images are downscaled to 224 × 224 pixels resolution before being used as inputs. The CNN-Control model performs consecutive CHF predictions on single-image inputs. This model was trained on a dataset of 4,368 images, half of which are taken during CHF and half of which are evenly distributed among non-CHF heat flux values.

CNN-Single. Another CNN using the same architecture and hyperparameters as CNN-Control is trained, but instead of conventional frame camera images, it is trained on images reconstructed from event data. Reconstructing event data into images involves accumulating events into time intervals and plotting them on a two-dimensional plane according to the event's polarity, which encodes whether the event recorded an increase or decrease in brightness. This model is designed to evaluate whether isolating the motion information recorded by events leads to an increase in accuracy or other performance metrics relative to CNN-Control. The training dataset is generated by reconstructing images from event streams as described previously, and 4,368 images are selected as in the case of CNN-Control. They are divided into categories according to the same description.

CNN-Square. Expanding upon the design of CNN-Single, another CNN is evaluated to determine how important temporal information is to the CHF detection problem. CNN-Square is designed to perform predictions on sequences of four reconstructed images (0.01 s duration), allowing for comparisons to be drawn with the performance of CNN-Control and CNN-Single. To minimize the data loss caused by image resolution adjustment, the four reconstructed images are pre-arranged into a 2 × 2 grid, which is then passed as input to the CNN-Square model. Hyperparameters and architecture remain unadjusted from CNN-Control and CNN-Single to isolate any differences in performance to the change in input. Since each input image to this model is composed of four separate images, the training dataset is approximately one-fourth the size of the datasets for CNN-Single and CNN-Square.

ConvLSTM. To better capture and store time-dependent relationships from sequential images, we train and test a convolutional long short-term memory (ConvLSTM) model, in which inputs consist of short videos as opposed to singular images. To mitigate the increase in computational burden, each image in the sequence is reshaped to a smaller resolution of 64 × 64, and 7 images (0.0175 s duration) are used for each input sequence. As with CNN-Square, the training dataset is derived from the same dataset as CNN-Control but is reduced by approximately a factor of 7 due to the grouping mechanism. Evaluation of this ConvLSTM provides insight as to whether improvements may be made to the performance of image-based prediction techniques by employing models specialized for predicting on spatiotemporal data.

eventLSTM. Finally, we demonstrate an eventLSTM model that forgoes the image and video mediums entirely, in favor of the more efficient as-is event streams with minimal pre-processing (Fig. 2(c)). The model input is a vector of 1,000 events, with each event containing 4 values corresponding to the event's (x, y, t, p). This dataset is obtained by converting video of the pool boiling experiments into events using *v2e*, and is preprocessed by segmenting the data into groups of 1,000

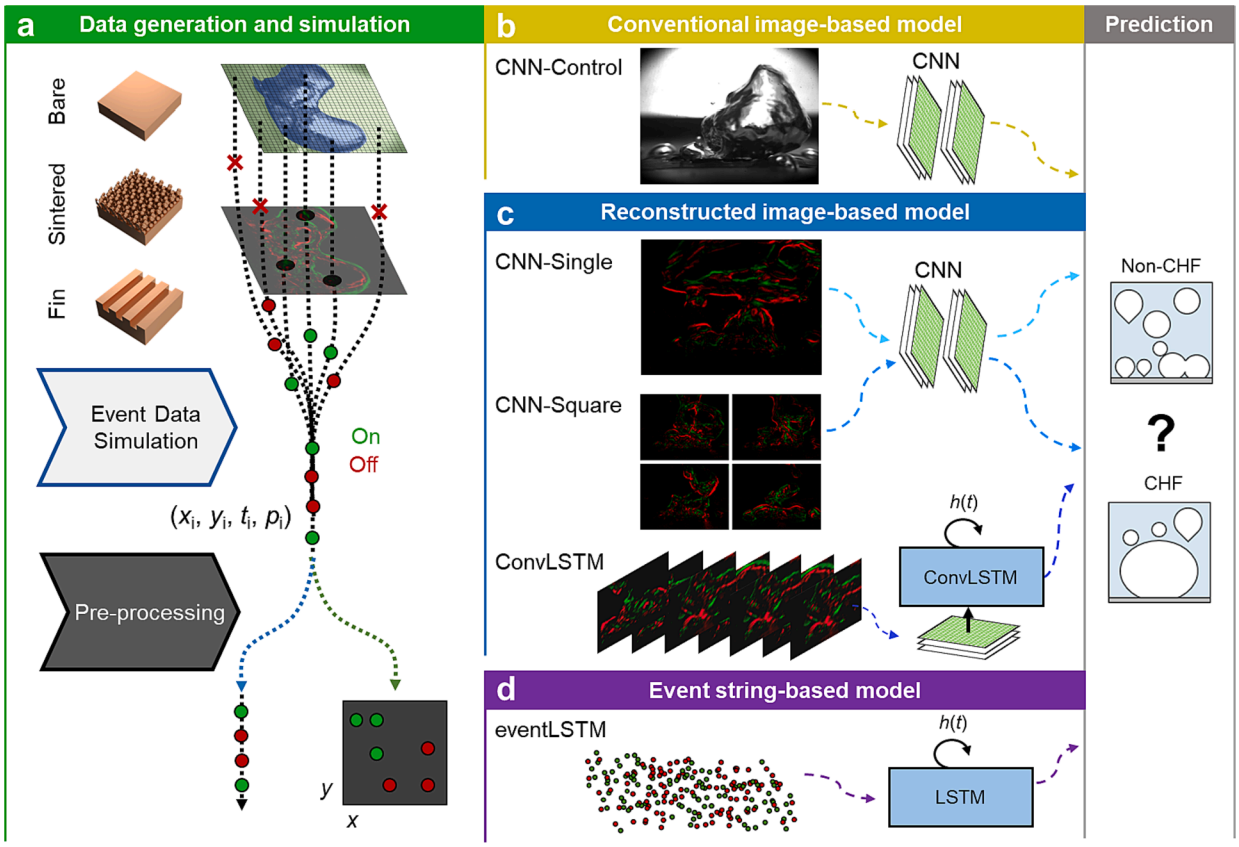

Fig. 2. Visualization of the framework used for this work. (a) Simulation of events using *v2e* which are then preprocessed either into an event-reconstructed image or an event stream segmented into intervals. (b) To establish both the performance and prediction time of a simple conventional image-based model, CNN-Control is trained on the original image data. (c) Three models are trained on event-reconstructed images, with CNN-Single, CNN-Square, and ConvLSTM preprocessing the images differently. (d) eventLSTM is trained on minimally preprocessed event data to exploit the differences between event and image data.

events. These are loaded in batches of 100, though both the batch size and the size of each event group are parameters that may be adjusted in future work for optimization. Since each event corresponds to a pixel-based brightness change (i.e., moving objects), this data is expected to allow for efficient and accurate predictions. Comparing the performance and prediction speed of the eventLSTM model to that of each previously discussed model, but particularly CNN-Control, is critical to understanding how well event data allows for circumventing the limitations of image-based CHF detection.

Each model is evaluated through three perspectives: effectiveness in the classification task, time taken for predictions, and ability to generalize for similar tasks. To explain these model performances, we use the term *classification performance* to denote how well a model classifies data by correctly distinguishing CHF from non-CHF data, the term *prediction time* to discuss the time taken for predictions, and the term *generalizability* for how well a model classifies CHF on datasets that do not match the dataset it was trained on.

Consistency is maintained by running all models on the same hardware, which would otherwise influence each model's overall performance and lead to biased comparisons. To establish a reference from which we can determine whether event-reconstructed images and event strings result in notable gains in overall performance, we assume that the CNN-Control model is a reasonably simplified representation of conventional visualization-based approaches.

2.3. Metrics

Ensuring that the trained models can effectively perform CHF detection with minimal error is imperative for practical applications. To evaluate our models, we utilize four conventional classification performance metrics called accuracy, precision, recall, and F-1 score.

Accuracy represents the proportion of all predictions that are correct. As the most common and easily interpretable metric of how well a model makes inferences, it is the primary indicator of performance. However, accuracy alone is an ineffective metric, especially when the class distribution in the data is uneven [24]. In such cases, a model that blindly classifies all data as the more common class will achieve relatively high accuracy despite not understanding the problem space. Thus, other metrics are also employed to provide a more robust understanding of how well each model performs on the data.

Precision represents the proportion of CHF classifications that are correct, or more generally the proportion of predictions of a particular class that are correct [24,25]. High precision indicates that we can trust a model's judgement when it labels a data point as CHF, while low precision suggests that a model is overzealous in its CHF predictions.

Recall is the proportion of CHF data points that the model correctly identifies, and in general terms is the proportion of a particular class of data points that are correctly predicted [24,25]. High recall indicates that the classifier can detect most actual CHF instances, whereas low recall means the model often fails to recognize CHF. Determining what constitutes an acceptable precision and recall is dependent on the problem context itself, as the two metrics are bound by a trade-off: increasing precision will require some amount of reduction in recall [25].

F1-score is a metric designed to account for both precision and recall, which is obtained by taking the harmonic mean of the two. The value of the F1-score is that it provides a single metric by which to evaluate a model that is more robust than accuracy. Furthermore, it serves as an easily interpretable metric that still accounts for the inverse relationship between precision and recall. The intuitive understanding of the F1-

score generally holds true: high values indicate better performance because they suggest that the model can discern positive cases without sacrificing the ability to isolate them from negatives.

These scores are based on TP, TN, FP, and FN, which respectively represent the numbers of true positives, true negatives, false positives, and false negatives in classifications. All metrics in this study are calculated by treating a particular data input as positive if it depicts CHF.

$$Accuracy = \frac{\sum(TP + TN)}{\sum(TP + FP + TN + FN)} \tag{1}$$

$$Precision = \frac{\sum TP}{\sum(TP + FP)} \tag{2}$$

$$Recall = \frac{\sum TP}{\sum(TP + FN)} \tag{3}$$

$$F1 - Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{4}$$

3. Results and discussion

3.1. Classification performance and prediction time

To determine the real-time viability of event-based CHF classification models, it is crucial to analyze the performance of each model both in the classification task and on its prediction time. While from a measurement perspective, the two are separate concepts, the speed of a model should ideally not come at a significant cost to classification performance. As shown in Fig. 3(a), the CNN-Control and eventLSTM models achieve nearly 100 % on all metrics while the other models show the metrics ranging between 92 and 99 %. These results indicate that models trained on event data can achieve similar classification performance to that of conventional image-based classifiers. For contexts such as CHF detection, where misclassification can have expensive and dangerous consequences, this similarity is a crucial prerequisite to justify substituting image-based classification with event-based classification. Achieving similar classification performance suggests that any improvements in prediction time do not come at the detriment of a model's understanding of the classification problem.

While the CNN-Control and eventLSTM models demonstrate similar strength with regards to classification performance metrics, a comparison of average model prediction times underlines the strength of event data. Each model's average prediction time and standard deviations are calculated and recorded, as shown in Fig. 3(b). The image-based models, including the ConvLSTM, average approximately 23 ms per prediction. In contrast, the eventLSTM achieves 7 ms per prediction, which is approximately a four-times increase. The data presented in Fig. 3(b) are averages of five separate trials where error bars represent one standard deviation. Notably, these prediction time values are recorded within a consistent Google Colab session, thus using identical hardware. By focusing GPU computing strategies, hardware selection and algorithm optimization, achieving a proportional speed-up is attainable for performing classification tasks. More broadly, the increase in speed under consistent conditions and equivalent classification performance suggests that a transition from conventional CNN-based models to event-based ones may generally lead to additional significant improvements in predictive speed, while maintaining equivalent classification capabilities.

3.2. Generalizability

When evaluating models, an additional point of consideration is how effectively a model generalizes to similar problems, such as detecting CHF in boiling experiments utilizing different surfaces. For this purpose, we evaluate how well eventLSTM models can generalize to other boiling surfaces when trained on data from varying sources. Several distributions of data are prepared in which three different pool boiling experimental conditions are represented in varying proportions, as seen in Fig. 4(a). A source dataset is compiled, containing data from potentially multiple boiling experiments, each of which is conducted on a different surface. This data is then split into training and validation datasets, allowing the model to learn from the training data while evaluating its performance on the validation data. A separate test dataset is prepared, containing only data from a single surface's boiling experiment. The test dataset is used to evaluate how well a model trained under a given data distribution performs for a specific target surface. Both the CNN-Control and eventLSTM models are trained on all distributions, each time being evaluated by their test accuracy on a dataset comprised entirely of data from Bare Cu boiling experiments. This emulates the scenario in which a model is trained for one data distribution but is generalized to predict on a different target distribution, in this case being purely Bare Cu.

The result of this evaluation is shown in Fig. 4(b), which also includes the performance of the CNN-Control model under the same conditions to serve as a comparison point. We observe that the eventLSTM model performs better than the CNN-Control model for data distributions composed of between 5 and 20 % Bare Cu training data. This indicates that as each model is provided less data from the target experimental conditions, the eventLSTM model maintains its classification performance better than the conventional CNN-Control model.

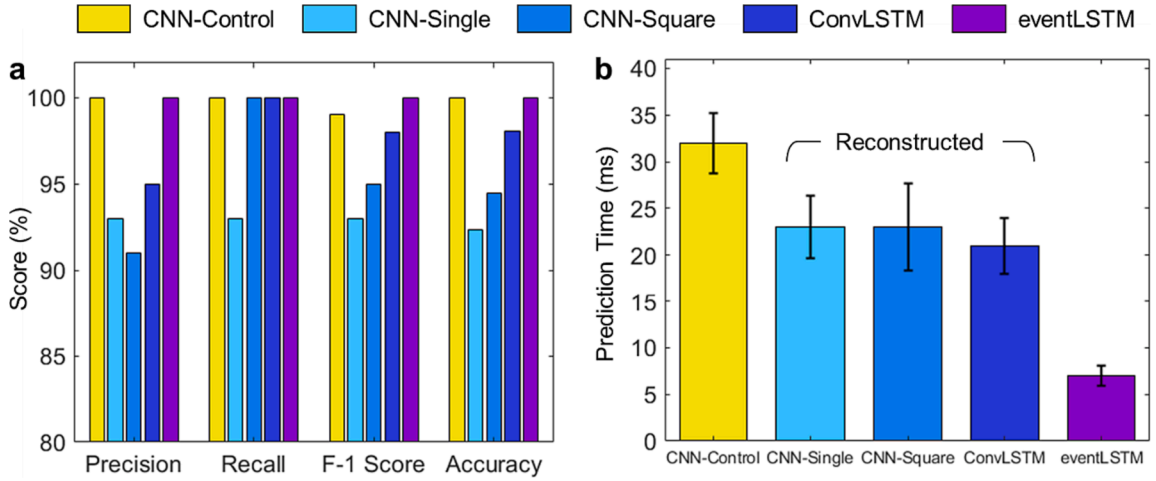

Fig. 3. (a) Graph depicting each model's classification performance metric values. (b) Milliseconds elapsed per prediction made by each model under testing circumstances.

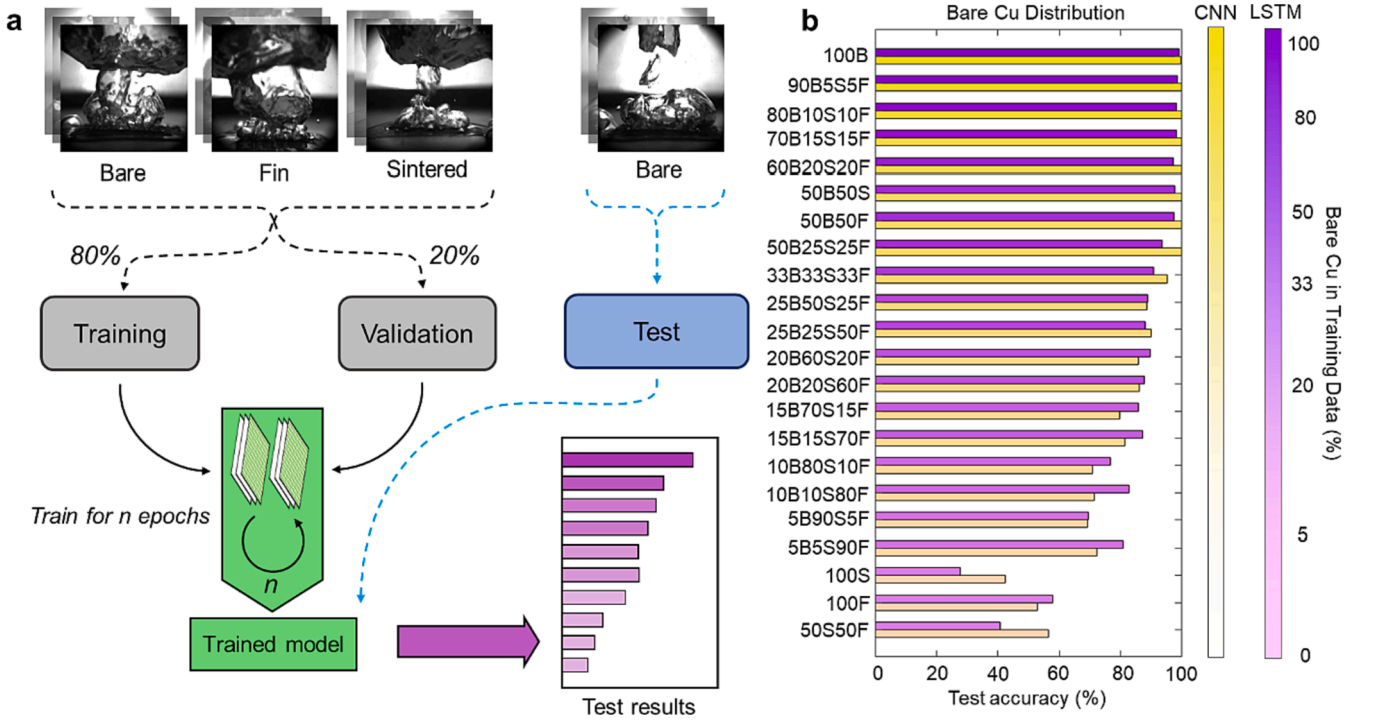

Fig. 4. (a) Process of preparing data for testing generalizability. (b) Graph depicting test accuracy of models under various training data distributions for the Bare Cu. Each distribution is labeled by the percentage of the dataset which contains each surface.

However, when the models are provided with no Bare Cu training data, the classification performance of the eventLSTM model drops sharply, allowing the CNN-Control model to outperform by approximately 19 % accuracy in some cases.

These results in model generalizability are reflective of how the spatio-temporal features learned by the eventLSTM model are powerful yet largely independent from one dataset to another, while the time-independent spatial features learned by the CNN-Control model encode more general information about the boiling process. As a model trained on time-series data, eventLSTM is designed to learn time-dependent boiling behavior, which varies significantly between different surface conditions. Our data suggests that eventLSTM models learn those time-dependent behaviors relatively well, even with limited data. However, when no Bare Cu data is provided, the model learns to predict features that are specific only to the remaining non-Bare Cu time-series data, leading to a drop in generalizability when scored on Bare Cu test accuracy. In contrast, the CNN-Control model makes predictions from single images, focusing on hierarchical visual patterns like the shape and size of bubbles. We speculate that the time-independent features are far more generalizable than sequential time-dependent features across different experimental conditions. Compared to the sequential motion of the boiling process, information like the shape of a bubble and its dominance within an image is far more generalizable between different experiments. Thus, in the extreme case of applying completely different training data to the Bare Cu CHF detection problem, the CNN-Control model appears to generalize better. However, once the models are trained on a small proportion of data from the target conditions, the eventLSTM model demonstrates its ability to maintain classification performance with fewer resources.

These results suggest that if limitations on collecting data are a concern, or if one intends to perform transfer learning with a relatively small amount of new training data, utilizing event-based models may offer additional advantages beyond those of prediction speed. Moreover, they invite further investigation into the potential benefits of neuromorphic imaging on generalizability, given the current findings solely pertain to surface modifications. Future analysis, particularly with non-simulated event data, may be required to determine how robust event-based models are to changes in diverse experimental conditions (e.g. camera placement, lighting conditions, and even working fluids).

4. Conclusion

The significant increase in prediction speed observed when transitioning from an image-based model to an event-based one has notable implications for the real-time applicability of CHF detection models. By achieving classification performance on par with conventional approaches in less time and through a power-efficient medium, developing real-time predictive models becomes a tangible possibility. Further extension of such work may lead to real-time forecasting of CHF, in which the time-series nature of event data is exploited to detect boiling crisis well in advance of its actual onset. If linked to a physical system, automatic maintenance of a particular targeted region within the heat transfer curve would become a much easier task. Optimization of event-based CHF detection models remains unexplored, but the current results indicate that such efforts are worth pursuing, as there are significant gains to be made in the transition of data formats. Furthermore, machine learning models that use image inputs for other motion-dependent contexts may be able to transition to event data to achieve similar reductions in prediction time. While future work is necessary to verify that the results identified in this work may be replicated with data generated by a physical neuromorphic camera, the framework demonstrated here may serve as a motivation and reference for developing real-time event-based models in various motion-based problem spaces.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Yoonjin Won reports financial support was provided by Office of Naval Research.

Data availability

Data will be made available on request.

Acknowledgments

The authors gratefully acknowledge the funding support from the Office of Naval Research (ONR) with Dr. Mark Spector as the program manager (Grant No. N00014-22-1-2063).

Appendix A. Supplementary material

See supplementary material for detailed discussion of performance metrics and model architectures. Supplementary data to this article can be found online at https://doi.org/10.1016/j.applthermaleng.2023.122004.

References

[1] Y. Suh, R. Bostanabad, Y. Won, Sci. Rep. 11 (2021) 5622.
[2] A.H. Howard, I. Mudawar, Int. J. Heat Mass Transf. 42 (9) (1999) 1665.
[3] S.M. Rassoulinejad-Mousavi, F. Al-Hindawi, T. Soori, A. Rokoni, H. Yoon, H. Hu, T. Wu, Y. Sun, Appl. Therm. Eng. 190 (2021), 116849.
[4] S.G. Kandlikar, J. Heat Transfer 123 (6) (2001) 1071.
[5] M.M. Rahman, E. Ölçeroglu, M. McCarthy, Langmuir 30 (37) (2014) 11225–11234.
[6] G.M. Hobold, A.K. da Silva, Int. J. Heat Mass Transf. 134 (2019) 511.
[7] G. Liang, I. Mudawar, Int. J. Heat Mass Transf. 117 (2018) 1352.
[8] J. Lee, A. Mohraz, Y. Won, Langmuir 38 (46) (2022) 14063.
[9] A. Rokoni, L. Zhang, T. Soori, H. Hu, T. Wu, Y. Sun, Int. J. Heat Mass Transf. 186 (2022), 122501.
[10] G.M. Hobold, A.K. da Silva, Int. J. Heat Mass Transf. 125 (2018) 1296.
[11] J. Lee, Y. Suh, M. Kuciej, P. Simadiris, M. Barako, Y. Won, Nanoscale 14 (2022) 13078.
[12] Y. Suh, S. Chang, P. Simadiris, T. Inouyet, M.J. Hoque, S. Khodakarami, C. Kharangate, N. Miljkovic, Y. Won, SSRN 4491956 (2023).
[13] S. Chang, Y. Suh, C. Shingote, C.N. Huang, I. Mudawar, C. Kharangate, Y. Won, SSRN 4458770 (2023).
[14] G.M. Hobold, A.K. da Silva, Int. J. Heat Mass Transf. 134 (2019) 262.
[15] R. Benosman, C. Clercq, X. Lagorce, S. Ieng, C. Bartolozzi, IEEE Trans. Neural Netw. Learn. Syst. 25 (2) (2013) 99.
[16] S.N. Gowda, C. Yuan, ACCV 2018 (2019) 581.
[17] Y. Wang, K. Lv, R. Huang, S. Song, L. Yang, G. Huang, Adv. Neural Inf. Process Syst. 33 (2020) 2432.
[18] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, D. Scaramuzza, IEEE Trans. Pattern Anal. Mach. Intell. 44 (1) (2022) 154.
[19] S. Yang, Y. Hou, Y. Shang, X. Zhong, AIP Adv. 12 (4) (2022), 045209.
[20] F. Demiray, J. Kim, Int. J. Heat Mass Transf. 47 (14–16) (2004) 3257.
[21] S. Durvasula, Y. Guan, N. Vijaykumar, IEEE Robot. Autom. Lett. 8 (6) (2023) 3174.
[22] Y. Hu, S.-C. Liu, T. Delbruck, in: 2021 Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. 1312 (2021).
[23] H. Jiang, D. Sun, V. Jampani, M.-H. Yang, E. Learned-Miller, J. Kautz, in: 2018 Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. 9000 (2018).
[24] B. Juba, H.S. Le, Proc. Innov. Appl. Artif. Intell. Conf. 33 (01) (2019) 4039.
[25] J. Davis, M. Goadrich, in: Proc. 23rd ICML, 233 (2006).