



Expert Monitoring: Human-Centered Concept Drift Detection in Machine Learning Operations

Joran Leest
 Vrije Universiteit
 Amsterdam
 The Netherlands
 j.g.leest@vu.nl

Claudia Raibulet
 Vrije Universiteit
 Amsterdam
 The Netherlands
 c.raibulet@vu.nl

Ilias Gerostathopoulos
 Vrije Universiteit
 Amsterdam
 The Netherlands
 i.g.gerostathopoulos@vu.nl

Patricia Lago
 Vrije Universiteit
 Amsterdam
 The Netherlands
 p.lago@vu.nl

ABSTRACT

We propose *Expert Monitoring*, an approach that leverages domain expertise to enhance the detection and mitigation of concept drift in machine learning (ML) models. Our approach supports practitioners by consolidating domain expertise related to concept drift-inducing events, making this expertise accessible to on-call personnel, and enabling automatic adaptability with expert oversight.

ACM Reference Format:

Joran Leest, Claudia Raibulet, Ilias Gerostathopoulos, and Patricia Lago. 2024. Expert Monitoring: Human-Centered Concept Drift Detection in Machine Learning Operations. In *New Ideas and Emerging Results (ICSE-NIER'24)*, April 14–20, 2024, Lisbon, Portugal. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3639476.3639771>

1 INTRODUCTION

The ubiquitous integration of machine learning (ML) in modern software systems marks a shift from deterministic behavior of software to behavior derived from stochastic processes. This transition brings a significant challenge: ensuring the consistent performance of ML models that are subject to changes in the data distribution, due to external events and data integrity issues [12]. This phenomenon, known as *data drift*, encompasses *feature drift* (changes in the input distribution, $P(X)$) and *concept drift* (changes in the conditional probability distribution of the target variable given an input, $P(Y|X)$) [29]. Notably, concept drift has a substantial influence on model performance and necessitates mitigative action [6, 14].

The growing use of ML underscores the importance of addressing concept drift, particularly in sensitive areas like credit card fraud detection, to prevent discriminatory actions [21]. We must scrutinize concept drift mitigation solutions, including insights from data mining [6, 17], and integrate them into the emerging ML operations (MLOps) framework used by practitioners [13].

In this paper, we explore MLOps practices and challenges, highlighting the inherent limitations of automated concept drift detection and mitigation, and motivate the the necessity of domain expertise (Section 2). Based on this understanding, we outline an approach to integrate expert knowledge into a monitoring system (Section 3). Finally, we discuss the key aspects of our approach (Section 4), and conclude by outlining our future plans (Section 5).



This work licensed under Creative Commons Attribution International 4.0 License.

ICSE-NIER'24, April 14–20, 2024, Lisbon, Portugal
 © 2024 Copyright held by the owner/author(s).
 ACM ISBN 979-8-4007-0500-7/24/04.
<https://doi.org/10.1145/3639476.3639771>

2 NAVIGATING THE DARKNESS: CONCEPT DRIFT IN PRACTICE

This section discusses the challenge of detecting and mitigating concept drift, and how domain expertise is utilized to address it.

2.1 The Practical Challenges of Concept Drift Detection – Where the Shadows Lie

We see the challenges in detecting and deciphering concept drift's latent aspects, often leaving practitioners "wandering in the dark".

2.1.1 Concept Drift Detection Without Labeled Data. Unlike feature drift, which can be readily observed in input data and its effects on model performance inferred through various estimation methods [3, 7, 10, 22, 23, 26], concept drift detection depends on monitoring metrics such as accuracy [6, 12]. This is challenging in real-world scenarios, as labeled data is often delayed or completely absent [12]. For example, in e-commerce churn prediction, churn determination only occurs after a specified time frame (e.g., a month or year) [1].

In the absence of labeled data, practitioners typically rely on detecting drift in the model's predictions and features, such as through the use of a Kolmogorov-Smirnov test, which can act as a proxy to infer the presence of concept drift and its effect on model performance [2, 12, 14]. This strategy makes sense because it leverages the common co-occurrence of feature drift and concept drift [6, 14, 16]. In real-world scenarios, events impact how data is generated, collected, and managed, leading to changes in our models that capture these processes [9]. Feature drift can act as a visible signal of these changes. Nonetheless, inferring concept drift from feature drift remains challenging, as their presence or severity does not always correlate [14]. As a result, triggering alerts for every instance of feature drift generates many false alarms [25–27].

To illustrate the challenge of concept drift detection without labeled data, let us consider a model that predicts whether a customer will churn based on the customer age and recent page visits (Fig. 1).

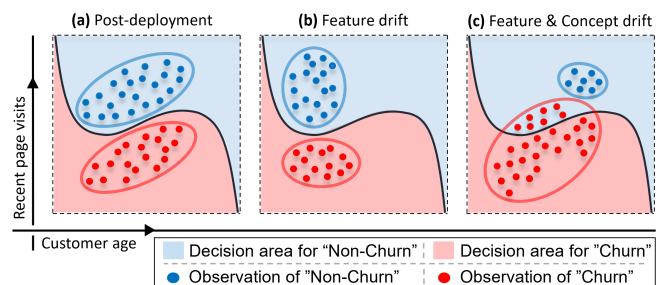


Figure 1: Data drift in a customer churn prediction model.

In Fig. 1, (a) depicts the post-deployment observations and the learned decision boundary (represented by the black line). After deployment, two events occur: (b) the web shop launches a marketing campaign targeting young people, causing drift in the customer age feature, and (c) a competitor’s marketing campaign for a new product line aimed at young people makes the web shop’s young customers switch to the competitor’s service, causing drift in the recent page visits feature. Event (b) does not affect customer satisfaction; the learned decision boundary remains valid. Conversely, event (c) negatively affects the (latent) customer satisfaction among the web shop’s younger customers, whom have become aware of the new offering. Consequently, concept drift occurs, rendering the learned mapping function invalid. The core problem emerges: **in the absence of labeled data, monitoring systems that rely on feature drift detection do not discern event (b) from (c).**

2.1.2 Deciphering the Nature of Concept drift Post-Detection. Even when the presence of concept drift can be confidently inferred, effective response selection requires an understanding of the detected drift’s characteristics [6, 11]. This includes the drift severity, recurrence, duration, and transition speed [29]. For example, in the case of recurrent drift, reactivating a previous model version might effectively resolve it [18], whereas abrupt drift might necessitate a complete model retrain. Conversely, for a short-lived drift duration, retraining the model is not desirable; instead, it might suffice to temporarily fall back on a more simple model.

Comprehending the nature of concept drift after its detection (e.g. whether it is a recurring event) remains a significant challenge [27], with current concept drift detection methods falling short in facilitating comprehension of the drift’s characteristics [17].

2.2 Domain Expertise – A Light in Dark Places, When All Other Lights Go Out

In recent works, Shankar et al. [25] and Shergadwala et al. [27] conducted insightful interview studies involving ML engineers. These studies highlighted the practical aspects of monitoring ML models and the strategies employed by ML engineers in practice. Their findings showed a common theme: **automated concept drift detection and mitigation is not a predominant tool in the arsenal of practitioners. Instead, human intervention and on-call rotations play a crucial role in monitoring ML models.** In doing so, practitioners are actively involved in maintaining the model’s overall and subgroup performances, aiming to optimize business value and ensure the model’s fairness, respectively [12, 25].

Now, let us first revisit the challenges we pinpointed in the previous section and consider how the domain expertise of a human-in-the-loop can address them. Afterwards, we examine the difficulties that arise when relying on domain expertise for model monitoring.

2.2.1 Addressing Concept Drift with Domain Expertise. To demonstrate how domain expertise can aid in detecting and mitigating concept drift, we again consider customer churn prediction and provide an illustrative example (see Fig. 2). Here, we present three illustrative instances of multivariate feature drift, along with an expert’s assessment of the potential underlying event and the presence of concept drift. Subsequently, we describe an appropriate response tailored to the nature of the detected drift.

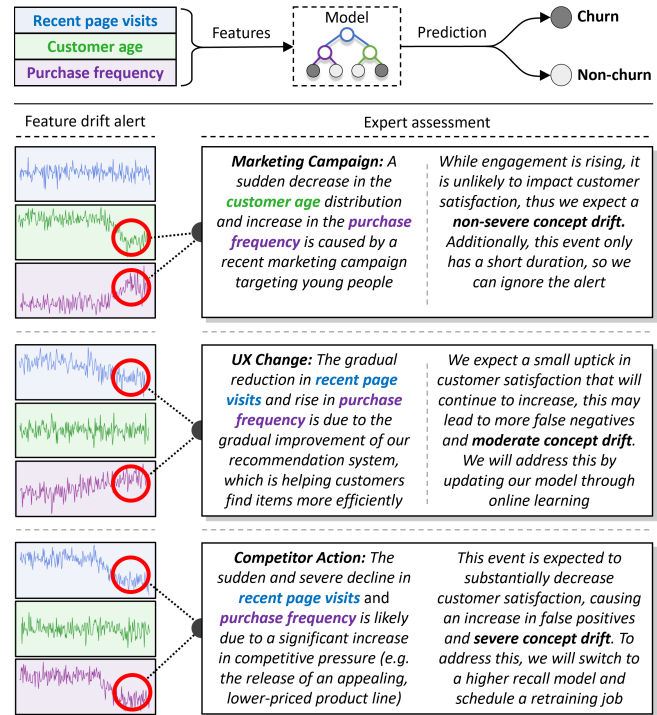


Figure 2: An illustrative case for customer churn prediction, showing expert assessments for three cases of feature drift.

The events depicted in Fig. 2 emphasize the vital importance of domain expertise in the process of concept drift detection and mitigation. While all events were effectively identified through feature drift detection, they varied significantly in their impact on model performance and the proper course of action.

2.2.2 Why Is It Hard to Rely on Domain Expertise? In the example (Fig. 2), we assumed the assessments were conducted by an expert with a deep understanding of the model and its application context. In practice, this responsibility often falls to on-call ML engineers, who oversee the monitoring of ML models [25, 27]. These models are often developed by different teams and operate in various contexts – processes that ML engineers have little to no involvement in! This proves indeed challenging, as quoted by an ML engineer: *"The pain point is dealing with that alert fatigue and the domain expertise necessary to know when to take action when on-call"* [25].

In addition to the need for domain expertise, ML engineers have also expressed a need for centralized model governance [27], as knowledge about ML models and their application context is often dispersed and inaccessible [20]. This issue becomes particularly difficult to manage for organizations that run numerous models, each with its unique feature set and context. For example, in addition to churn prediction, organizations might use models for tasks like demand prediction, product recommendation, and personalized search. The issue of acquiring and retaining domain expertise are exacerbated by factors such as staff turnover, limited documentation, and the need for extensive training [25, 27]. These challenges highlight the need to consolidate domain expertise and make it accessible to on-call ML engineers.

3 APPROACH

We propose a method called *Expert Monitoring* to address operational challenges. This method integrates domain expertise through *scenario specification*. It then makes this expertise accessible to on-call ML engineers via *scenario identification*, providing insights into the potential causes of feature drift upon detection (see Fig. 3).

3.1 Scenario Specification

We consolidate domain expertise through expert knowledge elicitation and retrospective analyses, creating a standardized resource for integration into our monitoring and response system.

3.1.1 Expert Knowledge Elicitation. Domain expertise is distributed among multiple experts in an organization, including marketing, product development, business strategy, and data engineering practitioners. We are interested in prior knowledge of events that can induce concept drift in the application context of an ML model. ML engineers collaborate with the domain experts to elicit scenarios of these events, using traditional requirements engineering methods such as interviews, focus groups, and observation studies [28].

3.1.2 Retrospective Analysis. In addition to the elicitation process, ML engineers, either in collaboration with domain experts or independently, conduct retrospective analyses. By examining the model’s historical performance and correlating performance drops with feature drift events (assuming access to labeled data), they can isolate recurring problematic events for future identification.

3.1.3 Scenario Specification Format. The acquired scenarios are compiled by the ML engineer and stored in a standardized format. Below, we provide a description of its components.

ML Model. The name and version of the ML model that is subject to concept drift in the specified scenario.

Scenario Description. This description provides the context for the event that can potentially induce concept drift.

Bayesian Model. We utilize Bayesian models to provide experts an intuitive method for incorporating their prior knowledge of how the feature distribution(s) would be affected under the specified scenarios. These models are central to the next phase of our approach, namely scenario detection, detailed in Section 3.2. They enable the estimation of the distributions of the relevant feature(s) as either univariate or multivariate, simultaneously quantifying the uncertainty in the experts’ subjective beliefs. Specifically, experts estimate the parameters (e.g. mean or standard deviation of an affected feature) as $\theta = Dist(location, spread)$, where the location is the estimated parameter value, and the spread indicates the expert’s uncertainty. A sharp distribution implies high certainty, while a wide one suggests low certainty. For example, for a highly certain estimate that the mean customer age in the marketing campaign scenario (Fig. 2) will be eighteen, we can define a normal distribution with location 18 and spread 1.

In addition to quantifying uncertainty, the representation of domain expertise can be flexibly determined in a fine-grained manner. Experts can: (1) select alternative distributions, such as a uniform distribution, to assign equal probabilities within specific ranges (e.g., ages 16 to 20); (2) provide relative estimates, in addition to absolute ones, in relation to the ML model’s training data (see Fig. 3); (3) define the distribution parameters of affected features for specific subgroups (e.g., in Fig. 1c, where the distribution of recent page visits can be estimated specifically for young people).

Scenario Understanding. This includes estimating concept drift characteristics like severity, transition speed, duration, and recurrence to inform response selection (see Section 3.2). Additionally, the likelihood of the scenario, whether it is common or rare, can also be specified based on prior knowledge, using a simple three-point scale (e.g., low, moderate, high) for consistency.

Scenario Response. An expert can optionally provide this response to guide the on-call ML engineer or automate scenario mitigation upon detection (see Section 3.2).

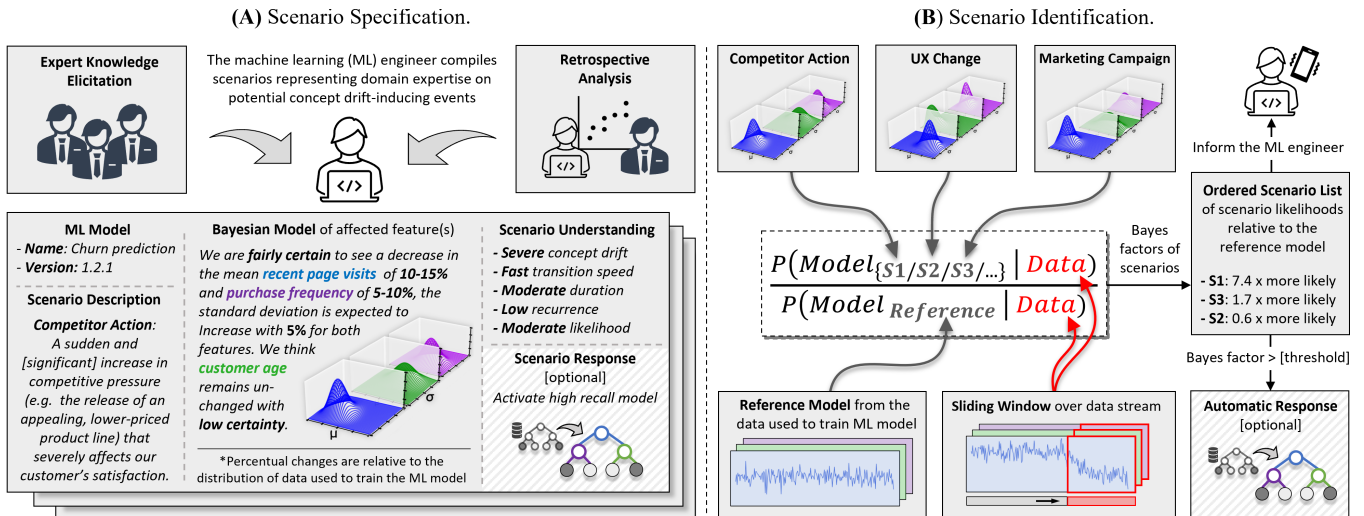


Figure 3: A visual depiction of our approach, *Expert Monitoring*. In step (A), ML engineers consolidate domain expertise within the organization using a standardized format. In step (B), upon detecting feature drift, scenarios are evaluated using Bayesian model comparison. Afterwards, the ML engineer is informed about potential causes, or an automated response is triggered.

3.2 Scenario Identification

At runtime, upon detecting feature drift, we infer the occurrence of a scenario using the Bayesian models defined in the prior step.

3.2.1 Bayesian Model Comparison. Each (Bayesian) scenario model is treated as a hypothesis and has its posterior probability $P(M|D)$ computed to assess its likelihood based on recent observations, obtained from a sliding window over the data stream. The posterior probability of a model is computed as follows:

$$P(M|D) = P(M)P(D|M) \quad (1)$$

Here, $P(M)$ represents the *scenario likelihood* of the model, reflecting the expert's belief about the likelihood of a scenario occurring, as discussed in Section 3.1.3. *Scenario likelihoods*, which can be set to equal by default, are specified on a three-point scale and normalized to sum to one. $P(D|M)$ denotes the *marginal likelihood* and is computed using the following integral:

$$P(D|M) = \int_1 \dots \int_n P(D|\theta, M)P(\theta|M) d\theta_1 \dots d\theta_n \quad (2)$$

This represents an n -dimensional integral over all parameters θ [8]. However, due to the impracticality of a closed-form solution, we instead use one of the following approximation methods: (1) Markov Chain Monte Carlo sampling [4], a computationally intensive method, or (2) calculating each feature's marginal likelihood in closed-form, by leveraging the conjugate prior assumption (the observed data and expert estimates follow the same distribution) [19], and then multiplying these likelihoods, assuming minimal inter-feature correlation.

After calculating the scenario models' posterior probabilities, we compare them with a reference model built from the observed parameter values in the ML model's training data. This comparison yields the Bayes factor [8], indicating the relative likelihood of each scenario model compared to the reference model.

In assessing the role of expert estimates in Bayesian model comparison for scenario identification, we find that the precision of these estimates directly influences accuracy (Fig. 4)¹. Specifically, scenarios identified using low-error (small deviation from the true parameter value) and high-certainty (low standard deviation specified in the estimate) estimates are typically more accurate. Moreover, even scenarios with higher estimate errors can be correctly identified if the associated uncertainty is correctly deemed high.




Figure 4: Detection accuracy vs. estimate uncertainty and error (in proportion relative to actual parameter values) on simulated scenarios, with a Bayes factor threshold of 5.

¹For more information, see https://github.com/JoranLeest/expert_monitoring

3.2.2 In Response to a Scenario. The Bayes factors from Bayesian model comparison are displayed with scenario identifiers, assisting ML engineers in response selection for feature drift. Our method also automates concept drift mitigation by triggering the top-ranked scenario's response based on a Bayes factor threshold.

4 DISCUSSION

In further developing our approach, there is an open question about extracting domain expertise on concept drift-inducing events through the reuse of requirement [28] and knowledge [24] elicitation methods. A second key question is: What is all the knowledge that we can incorporate regarding concept drift-inducing events? We have shown how a Bayesian model can be constructed to represent a scenario based on feature distributions, providing the base knowledge required from practitioners to explain feature drift occurrences. However, adopting the Bayesian framework offers the flexibility allows integrating extra knowledge and extending in various directions. For instance, experts might include scenario temporal distributions and likelihood at specific times, like higher sales in then summer. Furthermore, our method allows for updating expert estimates and facilitates the use of scenario-specific machine learning models, both of which help mitigate recurring scenarios.

The literature reveals a gap in understanding human-in-the-loop systems, especially domain expertise, for concept drift challenges. This area warrants more study to aid practitioners with appropriate processes, tools, and methodologies. While previous works offer methods for detecting feature drift-related model failures in practical settings [7, 10, 22, 23, 26], they overlook the latent issue of concept drift. More closely related to our research, Chen et al. [3] and Cobb et al. [5] incorporate domain expertise in performance estimation and feature drift detection, but also do not address concept drift. Our work uniquely integrates domain expertise in identifying and addressing concept drift, contributing to human-centered model monitoring [27]. We believe that scenario-based methods like ours, similar to those used in software architecture [15], are promising for advancing model monitoring.

5 FUTURE PLANS

We contend that the inherent intricacy of the human factors involved in our approach needs to be addressed through rigorous evaluation and collaboration with practitioners. Specifically, we identified the following research questions: (1) What domain expertise of concept drift-inducing events can be elicited, and represented with sufficient detail? (2) Can scenarios be identified through Bayesian model comparison with sufficient accuracy? (3) Is the *Expert Monitoring* approach perceived as helpful by ML engineers and does it enable them to improve on business-related metrics? To answer (1) and (2), we'll conduct action research via workshops and focus groups with diverse domain practitioners, refining our approach based on real-world industrial needs. For (3), we'll use surveys and interviews to gauge our approach's usefulness.

ACKNOWLEDGMENT

This research is supported by ExtremeXP, a project co-funded by the European Union Horizon Programme under Grant Agreement No. 101093164.

REFERENCES

- [1] Jaehyun Ahn, Junsik Hwang, Doyoung Kim, Hyukgeun Choi, and Shinjin Kang. 2020. A survey on churn analysis in various business domains. *IEEE Access* 8 (2020), 220816–220839.
- [2] Eric Breck, Shaoqing Cai, Eric Nielsen, Michael Salib, and D Sculley. 2017. The ML test score: A rubric for ML production readiness and technical debt reduction. In *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 1123–1132.
- [3] Mayee Chen, Karan Goel, Nimit S Sohoni, Fait Poms, Kayvon Fatahalian, and Christopher Ré. 2021. Mandoline: Model evaluation under distribution shift. In *International conference on machine learning*. PMLR, 1617–1629.
- [4] Siddhartha Chib. 2001. Markov chain Monte Carlo methods: computation and inference. *Handbook of econometrics* 5 (2001), 3569–3649.
- [5] Oliver Cobb and Arnaud Van Looveren. 2022. Context-aware drift detection. In *International Conference on Machine Learning*. PMLR, 4087–4111.
- [6] João Gama, Indrè Zliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM computing surveys (CSUR)* 46, 4 (2014), 1–37.
- [7] Saurabh Garg, Sivaraman Balakrishnan, Zachary C Lipton, Behnam Neyshabur, and Hanie Sedghi. 2022. Leveraging unlabeled data to predict out-of-distribution performance. *arXiv preprint arXiv:2201.04234* (2022).
- [8] Jeff Gill. 2014. *Bayesian methods: A social and behavioral sciences approach*. Vol. 20. CRC press.
- [9] Matthew Groh. 2022. Identifying the Context Shift between Test Benchmarks and Production Data. *arXiv preprint arXiv:2207.01059* (2022).
- [10] Devin Guillory, Vaishaal Shankar, Sayna Ebrahimi, Trevor Darrell, and Ludwig Schmidt. 2021. Predicting with confidence on unseen distributions. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1134–1144.
- [11] Meng Han, Zhiqiang Chen, Muhang Li, Hongxin Wu, and Xilong Zhang. 2022. A survey of active and passive concept drift handling methods. *Computational Intelligence* 38, 4 (2022), 1492–1535.
- [12] Chip Huyen. 2022. *Designing machine learning systems*. " O'Reilly Media, Inc."
- [13] Dominik Kreuzberger, Niklas Kühl, and Sebastian Hirschl. 2023. Machine learning operations (mlops): Overview, definition, and architecture. *IEEE Access* (2023).
- [14] Cloudera Fast Forward Labs. 2021. Inferring Concept Drift Without Labeled Data. (2021).
- [15] Joran Leest, Ilias Gerostathopoulos, and Claudia Raibulet. 2023. Evolvability of Machine Learning-based Systems: An Architectural Design Decision Framework. In *2023 IEEE 20th International Conference on Software Architecture Companion (ICSA-C)*. IEEE, 106–110.
- [16] Ziyi Liu, Rakshitha Godahewa, Kasun Bandara, and Christoph Bergmeir. 2023. Handling Concept Drift in Global Time Series Forecasting. *arXiv preprint arXiv:2304.01512* (2023).
- [17] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. 2019. Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering* 31, 12 (2019), 2346–2363. <https://doi.org/10.1109/TKDE.2018.2876857>
- [18] Ankur Mallick, Kevin Hsieh, Behnaz Arzani, and Gauri Joshi. 2022. Matchmaker: Data drift mitigation in machine learning for large-scale systems. *Proceedings of Machine Learning and Systems* 4 (2022), 77–94.
- [19] Kevin P Murphy. 2007. Conjugate Bayesian analysis of the Gaussian distribution. *def 1, 2σ* (2007), 16.
- [20] Felix Neutatz, Binger Chen, Ziawasch Abedjan, and Eugene Wu. 2021. From Cleaning before ML to Cleaning for ML. *IEEE Data Eng. Bull.* 44, 1 (2021), 24–41.
- [21] José Pombal, André F Cruz, João Bravo, Pedro Saleiro, Mário AT Figueiredo, and Pedro Bizarro. 2022. Understanding Unfairness in Fraud Detection through Model and Data Bias Interactions. *arXiv preprint arXiv:2207.06273* (2022).
- [22] Stephan Rabanser, Stephan Günemann, and Zachary Lipton. 2019. Failing loudly: An empirical study of methods for detecting dataset shift. *Advances in Neural Information Processing Systems* 32 (2019).
- [23] Tegjyot Singh Sethi and Mehmed Kantardzic. 2017. On the reliable detection of concept drift from streaming unlabeled data. *Expert Systems with Applications* 82 (2017), 77–99.
- [24] Nigel R Shadbolt, Paul R Smart, J Wilson, and S Sharples. 2015. Knowledge elicitation. *Evaluation of human work* (2015), 163–200.
- [25] Shreya Shankar, Rolando Garcia, Joseph M Hellerstein, and Aditya G Parameswaran. 2022. Operationalizing machine learning: An interview study. *arXiv preprint arXiv:2209.09125* (2022).
- [26] Shreya Shankar and Aditya Parameswaran. 2021. Towards observability for production machine learning pipelines. *arXiv preprint arXiv:2108.13557* (2021).
- [27] Murtuza N Shergadwala, Himabindu Lakkaraju, and Krishnaram Kenthapadi. 2022. A Human-Centric Take on Model Monitoring. *arXiv preprint arXiv:2206.02868* (2022).
- [28] Ian Sommerville. 2011. *Software engineering* (ed.). *America: Pearson Education Inc* (2011).
- [29] Geoffrey I Webb, Roy Hyde, Hong Cao, Hai Long Nguyen, and Francois Petitjean. 2016. Characterizing concept drift. *Data Mining and Knowledge Discovery* 30, 4 (2016), 964–994.