Towards a proof theory for reasoning about frequencies on series

Melissa Antonelli[1], Alessandro G. Buda[2], Leonardo Ceragioli[3] and Giuseppe Primiero[2]

[1]Universität Tübingen & CFvW, Doblerstraße 33, Tübingen, 72074, Germany
[2]University School for Advanced Studies IUSS Pavia, Piazza della Vittoria n. 15, Pavia, 27100, Italy
[3]University of Milan, Via Festa del Perdono 7, Milano, 20122, Italy

Abstract

The logic LTLFS was recently developed to reason directly about finite series of events. Its measure-sensitive operators support reasoning about observed frequencies and predictions of future occurrences, with potential applications in bias mitigation. However, LTLFS still lacks a formal proof system. We address this gap by introducing a labelled sequent calculus that uses labels associated with *measurable sets* of worlds and rules for counting. The calculus translates the quantitative semantics of LTLFS into syntactic rules for formal proofs.

Keywords

Temporal logic, Counting logic, Proof theory, Labelled sequent calculus, Frequencies of series

1. Introduction

Ensuring algorithmic fairness and preventing bias amplification increasingly relies on the real-time supervision and control of complex systems [1, 2, 3]. In this context, since retraining is costly, inference-only, post-hoc verification methods and guardrail architectures are crucial to intercept unacceptable outputs early, reducing unnecessary compute while ensuring trustworthy deployment [4, 5, 6, 7]. From the standpoint of formal analysis, a major related task consists in ensuring that an observed series of events, such as a sequence of outputs by a classifier, satisfies (quantitative) constraints over time. To address this, the temporal logic LTLFS [8] was recently introduced to express and verify properties of temporal experiments and comparing them with expected distributions of outputs, following the approach presented in [9, 10]. This system extends standard LTL semantics with original measure-sensitive modal operators to count event occurrences across different moments of a series, predict the probability of future outputs, and compare empirical and target distributions. Unlike traditional temporal logics [11, 12, 13], LTLFS enables us to evaluate past frequencies, with feasible extensions to formulate applications in bias mitigation [14].

This logic still lacks a formal proof system allowing us to syntactically investigate finite series and possibly providing ways to automate proof search. Given the quantitative nature of our logic,

BEWARE 2026 (Bias, Risk, Explainability and the role of Logic and Logic Programming), co-located with AIxIA 2026, October 8, 2026, Perugia, Italy.

melissa.antonelli@uni-tuebingen.de (M. Antonelli); alessandro.buda@iusspavia.it (A. G. Buda); leonardo.ceragioli@unimi.it (L. Ceragioli); giuseppe.primiero@unimi.it (G. Primiero)
0009-0006-9072-4847 (M. Antonelli); 0000-0002-5547-7514 (A. G. Buda); 0000-0001-5250-9720 (L. Ceragioli); 0000-0003-3264-7100 (G. Primiero)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR-WS.org/Vol-4332/paper9.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-17

this task is not trivial. Here, we present the very first core ideas to overcome the main structural challenges and formulate a sound and complete deductive system for LTLFS. Overall, our approach generalizes the standard framework of labelled sequent calculi (remarkably effective for modal [15, 16] and temporal logics [17]), as well as systems designed to accommodate causal relations among graphs representing features in classification tasks [18, 19], by introducing more general labels associated with *measurable sets* of worlds and inference rules capable of "counting".[1]

2. A gentle introduction to LTLFS

In this section, we summarize the formal language and semantics of a simplified version of LTLFS, called LTLFS_0, in which, for brevity, standard LTL operators are not considered and nested modalities are not allowed. We assume familiarity with basic measure theory (see [20]), required to define the semantics of LTLFS. The language of LTLFS_0 is a standard propositional language, endowed with five original operators. In what follows, we will use $p_1, p_2, \ldots$ for atomic propositions and $A, B, \ldots$ for possibly molecular ones. We use $\mathfrak{P}$ to denote the set of atomic variables.

Definition 1 (Language of LTLFS_0)**.** Formulas of LTLFS_0 are defined by the grammar below:

$$A ::= p \mid \neg A \mid A \wedge A \mid A \vee A \mid \Box_q B \mid \blacksquare_q B \mid \circledcirc_q B \mid \triangleright_q B \mid \star_q B$$

where $p \in \mathfrak{P}$, $q \in \mathbb{Q}_{[0,1]}$ and B is a classical propositional formula.

With the language $\mathcal{L}_{\mathsf{LTLFS}_0}$ one can naturally expresses properties of series of events and reason on output frequencies. The intuitive meaning of the operators is as follows:

- $\Box_q B$: until time t_m, the event corresponding to B has been observed *at least* $q \in \mathbb{Q}_{[0,1]}$ of the times;
- $\blacksquare_q B$: according to a given target frequency, up until time t_m, the event corresponding to B could have happened *at least* q of the times;
- $\circledcirc_q B$: until time t_m the event corresponding to B has been obtained at least $q \times n$ times, where n is the total number of trials executed in the current series;
- $\triangleright_q B$: assuming the current realization is associated with a given distribution, the probability that the event corresponding to B happens at the next time t_{m+1} is at least q;
- $\star_q B$: according to a given probability distribution associated with our series of experiments, the event corresponding to B has a global probability of at least q.

Example 1. Consider a four-toss coin experiment with outcome $HHTH$ (H for heads, T for tails). Due to the operator $\circledcirc$, we can track the number of occurrences relative to the total number of trials in the entire series (here, $n = 4$). For example, the formula $\circledcirc_{1/4} p_H$ holds at t_1, $\circledcirc_{2/4} p_H$ at both t_2 and t_3, and $\circledcirc_{3/4} p_H$ at t_4, formally expressing that at time t_1, t_2, t_3, and t_4, the coin returned heads, at least 1, 2, 2, and 3 times, respectively.

[1] For brevity, this presentation is informal and as compact as possible. Full details, including formal definitions and soundness proofs, can be found in the eponymous extended version.

Semantically, [8] introduces a Kripke-style semantics defined over models made of a *finite* set of worlds interpreted as time instants, a special accessibility relation to be interpreted as a linear time series, and a set of assignments corresponding to a well-defined frequency distribution. Worlds are used to model instants of time at which events occur. Although the semantics is standard in the sense that formulas are either true or false in a given model and interpretation, boxed formulas are evaluated considering *sets* of worlds and assignments. Intuitively, each world represents one position in a finite sequence of experimental outcomes, and each assignment labels every world with exactly one outcome atom. The assignments in the reference distribution represent sequences whose overall proportions agree with the reference probability measure, whereas the outcomes actually observed are stored in an appropriate variable decorating the consequence relation. The order among events lets us consider the initial segment $\{w_1, \ldots, w_m\}$ at w_m, where the current series can be evaluated. The notion of validity in a model is defined in the standard way. For complete semantics and details on the world-language measure relationship, see [8].

3. Towards a proof theory for LTLFS

In this section, we introduce preliminary ideas to develop a sound and complete sequent calculus for LTLFS. We consider here only the toy fragment LTLFS_0 to avoid technical issues related to, for example, the nesting of modalities, while providing a clear intuition of the core ingredients needed to obtain sound (yet expressive) rules for selected case-study operators. The idea is to follow the well-established approach of [16, 15], internalizing the modal semantics into the syntax of the calculus. Here, this methodology is generalized to our non-standard modal operators by considering labels corresponding to *measurable sets* of worlds. This provides our formal mechanism for counting worlds and expressing probabilities.

For the sake of brevity, in the current introductory presentation, we confine our analysis to one operator only. However, the core intuition of combining rules for relational properties [16, 17] with relational atoms that express properties of *sets* of worlds scales to all our operators. This provides the basis for developing a *full*, sound and complete calculus for LTLFS, which we leave as future work.

Notation 1. Let $k, o, \cdots \in \mathcal{W}$ be variables for worlds, $\alpha, \beta, \cdots \in \mathcal{S}$ for sets, and $t, s, \cdots \in \mathcal{W} \cup \mathcal{S}$ for terms. For readability, we assume that whenever $k_{i_1}, \ldots, k_{i_m}$ are listed, the worlds are ordered according to their subscripts (i.e., $k_{i_1} \leq k_{i_2}$, and so on). Additionally, below in rules R^2_μ, $\mathsf{L}\circledcirc$, $\mathsf{R}\circledcirc$, we assume that $k_{i_1}, \ldots, k_{i_m}$ correspond to *different* worlds (i.e., for any $i_j, i_{j'}$, the formula $k_{i_j} \neq k_{i_{j'}}$ occurs in the left part of the sequent).

Definition 2. Let $* \in \{\geq, >, \leq, <, =, \neq\}$. The language of the calculus is made of relational atoms and labelled formulas. *Relational atoms* are the expressions below:

- $s * t$, expresses that a world or set s is in the corresponding relation with t;
- $k \in \alpha$ expresses that $w = [\![k]\!]$ is a world of the set $[\![\alpha]\!]$;
- $\#(\alpha) * m$ expresses that the number of worlds in α is in relation $*$ with $m \in \mathbb{N}$;

For any $A \in \mathcal{L}_{\mathsf{LTLFS}_0}$, the *labelled formulas* $k : A$ and $\alpha : A$ intuitively express that the formula A is true in world k and in *all* worlds z such that $z \in \alpha$, resp.

The semantics for relational atoms and labelled formulas are defined as expected, consistently with the intuitive meaning of $*$ and set inclusion. Specifically, relational atoms provide information regarding the structure of the model and the number of worlds in the set.

Definition 3. Sequents are expressions of the form:

$$\Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta$$

where Γ and Δ are multi-sets of labelled formulas and relational atoms, $n \in \mathbb{N}$ corresponds to the cardinality of the model, and Φ_Σ is a set of assignments compatible with a distribution associated with the atoms occurring in Γ and Δ.

Sequents are interpreted in the standard way [15], additionally expressing that the total number of worlds is equal to $n \in \mathbb{N}$ and that valuations respect a given distribution Φ_Σ.

The rules defining our labelled calculus are expressed in the labelled language and include all expected (generalized) initial sequents, standard logical rules for connectives in modal systems (see [15]), and rules for linear order (analogously to those introduced for temporal logics, see e.g. [17]). To these, we add four initial sequents, two rules characterizing probability distribution, and two logical rules for each modal operator. We call the resulting calculus $\mathbf{GLTLFS_0}$, see Fig. 1. The intuitive meaning of these rules is as expected. For example, the rule $\mathsf{L}\circledcirc$ expresses how a frequency assumption can be used in a derivation. Bottom-up, assuming A holds with frequency $\frac{m}{n}$ in the world corresponding to k_l allows us to extract m different positions $k_{i_1}, \ldots, k_{i_m}$ none later than k_l, in which A holds. Top-down, if these occurrences, together with Γ, suffice to derive Δ, then the frequency assumption $k_l : \circledcirc_{\frac{m}{n}} A$, together with Γ, also suffices. For this reading, $k_{i_j} \leq k_l$ should hold for every $j = 1, \ldots, m$, the positions should be distinct and belong to the relevant n-position interval. The rule $\mathsf{R}\circledcirc$ introduces a frequency claim: if, under the same background assumptions, A can be established at m distinct positions no later than k_l, then we may conclude at k_l that A has frequency $\frac{m}{n}$. The constraints in D ensure that each occurrence is counted only once, while n is the size of the relevant observation interval. Strictly speaking, these premises establish at least m occurrences.

As desirable, the calculus is sound with respect to the semantics of LTLFS_0.

Proposition 1. *If* $\Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta$ *is derivable in* $\mathbf{GLTLFS_0}$*, then it is valid.*

Proof sketch. By induction on the derivation height. All cases except for those corresponding to measure and modal rules are standard. We consider $\circledcirc$ as an example. By IH the premiss is valid. Suppose, by contradiction, that $[\![\cdot]\!]$ does not satisfy the conclusion. That is, assume $[\![\cdot]\!]$ satisfies all formulas in Γ and $k_l : \circledcirc_{\frac{m}{n}} F$, and no formula in Δ. W.l.o.g., let $[\![k_l]\!] = w_l$. By the semantics of $\circledcirc$, the set $S = \{w_r : r \leq l \wedge \mathbb{M}, w_r \vDash F\}$ satisfies $\frac{|S|}{n} \geq \frac{m}{n}$, i.e., $|S| \geq m$. Choose m distinct worlds $v_{j_1}, \ldots, v_{j_m} \in S$, ordered as in Notation 1, and let $[\![\cdot]\!]' = [\![\cdot]\!][k_{i_1} \mapsto v_{j_1}, \ldots, k_{i_m} \mapsto v_{j_m}]$. Since $k_{i_1}, \ldots, k_{i_m}$ are fresh, they do not occur in Γ, Δ or $k_l : \circledcirc_{\frac{m}{n}} F$. Hence, $[\![\cdot]\!]'$ agrees with $[\![\cdot]\!]$ on all labels occurring there: it satisfies all formulas in Γ, no formula in Δ, and $[\![k_l]\!]' = w_l$. Moreover, for every $s \in \{1, \ldots, m\}$: $[\![k_{i_s}]\!]' = v_{j_s} \leq w_l$ and $[\![k_{i_s}]\!]' \vDash F$, because $v_{j_s} \in S$; the v_{j_s} are pairwise distinct and ordered. Thus, $[\![\cdot]\!]'$ satisfies the whole antecedent of the premiss and no formula in Δ, contradicting the validity of the premiss. Hence, the conclusion is valid. $\square$

Initial sequents

$$s * t, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, s * t \qquad k_1, \ldots, k_n \in \alpha, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, k_1, \ldots, k_n \in \alpha$$
$$\#(\alpha) * m, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, \#(\alpha) * m \qquad t : p, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, t : p$$

Sample of logical and relational rules (see [15, 17])

$$\frac{k : A, k : B, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta}{k : A \wedge B, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta} \mathsf{L}\wedge \qquad \frac{k_1 \leq k_3, k_1 \leq k_2, k_2 \leq k_3, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta}{k_1 \leq k_2, k_2 \leq k_3, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta} \mathsf{Tr}$$

Measure rules

$$\frac{\#(\alpha) = m, \alpha : p, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta}{\Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta} \mathsf{R}^1_\mu \qquad \frac{\#(\alpha) = m, \alpha : A, k_{i_1} : A, \ldots k_{i_m} : A, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta}{\#(\alpha) = m, \alpha : A, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta} \mathsf{R}^2_\mu$$

where $\mu_{\Phi_\Sigma}([\![p]\!]) = \frac{m}{n}$ and α is fresh

where $k_{i_1}, \ldots k_{i_m}$ are fresh and *distinct* variables

Modal rules

$$\frac{k_{i_j} \leq k_l, k_{i_1} : A, \ldots, k_{i_m} : A, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta}{k_l : \circledcirc_{\frac{m}{n}} A, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta} \mathsf{L}\circledcirc$$

with $j \in \{1, \ldots, m\}$ and $k_{i_1}, \ldots, k_{i_m}$ are fresh and *distinct*

$$\frac{D, k_{i_1} \leq k_l, \ldots, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, k_{i_1} : A \qquad \ldots \qquad D, \ldots, k_{i_m} \leq k_l, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, k_{i_m} : A}{D, k_{i_1} \leq k_l, \ldots, k_{i_m} \leq k_l, \Gamma \Rightarrow^{\Phi_\Sigma, n} \Delta, k_l : \circledcirc_{\frac{m}{n}} A} \mathsf{R}\circledcirc$$

being $D = \{k_{i_j} \neq k_{i_{j'}} \ : \ 1 \leq j < j' \leq m\}$

Figure 1: Rules of $\mathbf{GLTLFS_0}$

4. Preliminary conclusion and future work

In this work, we have formulated initial steps toward developing a labelled sequent calculus based on original set labels and new measuring rules. Future work will extend this calculus to the full language of LTLFS, considering all modal operators and allowing nested modalities. This will enable us to analyze decision complexity, investigate rule invertibility and (automated) proof search, and further explore the non-standard distinction between global and local (world-relative) validity. It will also clarify the connection between our logic and existing probabilistic extensions of LTL, counting logics, and branching presentations (see, e.g., [14]). Beyond LTLFS, we believe that this framework, endowed with labelled sets, could offer a natural, general foundation for other proof systems in quantitative settings.

Declaration on Generative AI

During the preparation of this work, the authors used Claude to proof read the paper. The authors take full responsibility for the publication's content.

References

[1] P. Seshadri, S. Singh, Y. Elazar, The bias amplification paradox in text-to-image generation, 2023. URL: https://arxiv.org/abs/2308.00755.

[2] A. G. Buda, G. Coraglia, F. A. Genco, C. Manganini, G. Primiero, Bias amplification chains in ML-based systems with an application to credit scoring, in: Proceedings of the International Workshop on Bias, Risk, Explainability and the role of Logic and Logic Programming (BEWARE 2023), CEUR-WS, 2023.

[3] Y. Hirota, Y. Nakashima, N. Garcia, Quantifying societal bias amplification in image captioning, in: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022, pp. 13440–13449.

[4] Y. Guo, M. Guo, J. Su, Z. Yang, M. Zhu, Bias in large language models: Origin, evaluation, and mitigation, 2024. URL: https://arxiv.org/abs/2411.10915.

[5] S. Han, S. Avestimehr, C. He, Bridging the safety gap: A guardrail pipeline for trustworthy llm inferences, 2025. arXiv:2502.08142.

[6] H. Zhou, Z. Feng, Z. Zhu, J. Qian, K. Mao, Unibias: unveiling and mitigating LLM bias through internal attention and ffn manipulation, in: Proceedings of the 38th International Conference on Neural Information Processing Systems (NIPS 2024), Curran Associates Inc., Red Hook, NY, USA, 2024.

[7] X. Cheng, R. Chen, H. Zan, Y. Jia, M. Peng, Biasfilter: An inference-time debiasing framework for large language models, 2025. URL: https://arxiv.org/abs/2505.23829.

[8] M. Antonelli, L. Ceragioli, A. Buda, G. Primiero, A linear temporal logic of frequencies on series of events, 2026. URL: https://arxiv.org/abs/2604.10669.

[9] E. Kubyshkina, G. Primiero, A possible worlds semantics for trustworthy non-deterministic computations, International Journal of Approximate Reasoning 172 (2024) 109212.

[10] F. A. D'Asaro, F. A. Genco, G. Primiero, Checking trustworthiness of probabilistic computations in a typed natural deduction system, Journal of Logic and Computation 35 (2025) exaf003.

[11] B. Finkbeiner, F. Klein, K. H. Rath, Counting models of linear-time temporal logic, in: Language and Automata Theory and Applications - 8th International Conference (LATA 2014), 2014, pp. 304–315.

[12] H. Hansson, B. Jonsson, A logic for reasoning about time and reliability, Formal Aspects of Computation 6 (1994) 512–535.

[13] F. Laroussinie, A. Meyer, E. Petonnet, Counting LTL, in: 17th International Symposium on Temporal Representation and Reasoning, IEEE Computer Society, Paris, France, 2010, pp. 51–58.

[14] A. G. Buda, G. Primiero, L. Ceragioli, M. Antonelli, Counting worlds branching time semantics for post-hoc bias mitigation in generative AI, in: Proceedings of the 8th International Workshop on Artificial Intelligence and Formal Verification, Logic, Automata, and Synthesis (OVERLAY 2026), 2026.

[15] S. Negri, J. von Plato, Proof Analysis: A Contribution to Hilbert's Last Problem, Cambridge University Press, 2011.

[16] S. Negri, Proof analysis in modal logic, Journal of Philosophical Logic 34 (2005) 507–544.

[17] B. Boretti, S. Negri, Decidability for priorean linear time using a fixed-point labelled

calculus, in: M. Giese, A. Waaler (Eds.), Automated Reasoning with Analytic Tableaux and Related Methods, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 108–122.
[18] L. Ceragioli, G. Primiero, A proof system with causal labels (part i): checking individual fairness and intersectionality, in: Proceedings of the 7th International Workshop on Artificial Intelligence and Formal Verification, Logic, Automata, and Synthesis (OVERLAY 2025), 2025, pp. 173–179.
[19] L. Ceragioli, G. Primiero, A proof system with causal labels (part ii): checking counterfactual fairness, in: Proceedings of the 7th International Workshop on Artificial Intelligence and Formal Verification, Logic, Automata, and Synthesis (OVERLAY 2025), 2025, pp. 181–187.
[20] P. Billingsley, Probability and Measure, John Wiley and sons, 1995.