Enhancing the LTI Protocol to Support Alternative Grading Practices in Online Learning Tools

Ilse van Duijl, Virginia Polytechnic Institute and State University
Dr. Bob Edmison, Virginia Polytechnic Institute and State University

Dr. Edmison is a Collegiate Associate Professor in the Department of Computer Science at Virginia Tech. His research focuses on alternative grading practices, accessibility, and the tools to support them.

Prof. Stephen H Edwards, Virginia Polytechnic Institute and State University

Stephen H. Edwards is a Professor and the Associate Department Head for Undergraduate Studies in the Department of Computer Science at Virginia Tech, where he has been teaching since 1996. He received his B.S. in electrical engineering from Caltech, and M.S. and Ph.D. in computer and information science from The Ohio State University. His research interests include computer science education, software testing, software engineering, and programming languages. He is the project lead for Web-CAT, the most widely used open-source automated grading system in the world. Web-CAT is known for allowing instructors to grade students based on how well they test their own code. In addition, his research group has produced a number of other open-source tools used in classrooms at many other institutions. Currently, he is researching innovative methods for giving feedback to students as they work on assignments to provide a more welcoming experience for students, recognizing the effort they put in and the accomplishments they make as they work on solutions. The goals of his research are to strengthen growth mindset beliefs while encouraging deliberate practice, self-checking, and skill improvement as students work.

Dr. Adrienne Decker, University at Buffalo, The State University of New York

Adrienne Decker is a faculty member in the newly formed Department of Engineering Education at the University at Buffalo. She has been studying computing education and teaching for over 15 years, and is interested in broadening participation, evaluating t

Audrey Rorrer, University of North Carolina at Charlotte

Audrey Rorrer, PhD, is a Research Associate Professor in the Computer Science Department at UNC Charlotte, where she also serves as Assistant Director of the Center for Education Innovation & Research. Dr. Rorrer's scholarship areas include the science of broadening participation in computing, SoBP, which is a recognized domain of critical importance in STEM workforce development and educational programming. Her work has focused on educational programs, outreach and collective impact activities that expand the national pipeline into STEM careers. College student development and Faculty career development are central themes across her body of work.

Dr. Manuel A. Pérez-Quiñones, University of North Carolina at Charlotte

Dr. Manuel A. Pérez Quiñones is a Professor of Software and Information Systems at UNC at Charlotte. His research interests include diversity issues in computing, CS education, and human-computer interaction. He currently serves on the Committee on Women in Science, Engineering and Medicine at the National Academies and served as Program Officer at the National Science Foundation. His efforts to diversify computing have been recognized with an ACM Distinguished Member status (2019); the CRA A. Nico Habermann award (2018); and the Richard A. Tapia Achievement Award (2017). He is originally from San Juan, Puerto Rico.

Enhancing the LTI Protocol to Support Alternative Grading Practices in Online Learning Tools

1 Abstract

Educators are increasingly adopting alternative grading practices designed to increase student engagement, put greater focus on learning outcomes instead of points, and ensure all students have access to the same learning opportunities. Such practices often include performance-based grading scales that do not use points, flexible deadlines, resubmission of low-scoring work, and clear, actionable feedback. Unfortunately, existing online educational tools often focus on points-based grading only, with numeric scores being the primary feedback mechanism. The Learning Tools Interoperability (LTI) protocol is a key communications protocol between learning management systems and other online educational tools that embodies the points-based focus of tool builders. Lack of tool support for alternative grading practices is a significant impediment to the widespread adoption of these practices. This paper introduces the EGP broker, an innovative, intermediary tool designed to use a proxying strategy to inject sophisticated support for new grading schemes and enhanced feedback into the LTI protocol by operating between various online learning tools and the corresponding learning management system, effectively bridging the technological gap that currently hinders pedagogical progress. The primary goal of the EGP broker is to serve as a universal translation layer, normalizing data and communication between disparate educational applications to facilitate alternative grading workflows. The tool is implemented as a robust web services application built on a Nuxt 4 and Nitro architecture with a PostgreSQL backend. It supports both LTI 1.3 and legacy LTI 1.1 applications to facilitate cross-protocol communication. The system's design is highly modular, allowing for complex grade transformations such as converting multiple tool attempts into a single categorical mastery score. This paper describes the limits with the existing LTI protocol with respect to newer alternative grading practices, explains how a middleware approach can help bridge the gap, presents the design of the EGP broker itself, and describes early experiences with the approach using scenarios based on non-points-based grading, flexible deadlines, resubmission of work, and enhanced feedback.

2 Introduction

The foundation of course evaluation in computer science education has traditionally relied on a points-based, weighted-average grading system. While widespread, these traditional grading practices are increasingly recognized as exacerbating systemic inequity that can disproportionately and negatively impact students from marginalized communities [1]. This

recognition has fueled a growing movement toward Equitable Grading Practices (EGP), where the focus shifts from the accumulation of points to the accurate assessment of demonstrable learning mastery [2]. The widespread adoption of these practices in computing education is currently impeded by a "technological momentum" that favors traditional, points-based systems over more nuanced assessment models [3, 4]. EGP models, such as specifications grading and mastery grading, require instructors to adopt non-point-based, categorical scales (e.g., complete/incomplete, mastery/developing) that better reflect student competence.

The core challenge to adopting these innovative educational practices is a fundamental misalignment between the desired instructional methodology and the capabilities of the prevailing technology stack. While instructors rely on external learning tools for specialized, automated assessment, the grades from these systems must ultimately be recorded in the central gradebook of a Learning Management System (LMS), such as Canvas[1] or Moodle[2]. The communication between these platforms is governed by the Learning Tools Interoperability (LTI) standard [5]. Crucially, this LTI-mandated grade passback mechanism is built for the transfer of a single, normalized numerical score. This forces instructors to either abandon their flexible grading schemes or manually translate rich categorical data into a simplified, imprecise, numerical score that the LMS will accept.

This paper addresses this critical technology gap by presenting the design and implementation of a specialized LTI Proxy. The proxy serves as essential middleware, positioned between the external tool and the LMS, to facilitate dynamic grade translation. It intercepts non-standard, categorical grade data from the external tool, applies instructor-defined logic to convert that data into an LTI-compliant numerical score, and securely passes the result to the LMS gradebook. This intervention provides the necessary technical flexibility to decouple the instructor's assessment scheme from the LMS's rigid data structure, thereby making it practical and sustainable to implement innovative and equitable grading practices in high-enrollment CS courses.

3 Related Work

The shift toward Equitable Grading Practices (EGP) in computer science education is driven by a documented tension between traditional, points-based grading and student success. To address the problem of technical inflexibility that hinders the adoption of EGP, this paper first establishes the educational need for alternative grading and then details the technical constraints imposed by current Learning Tools Interoperability (LTI) standards and Learning Management Systems (LMS).

3.1 Pedagogical Foundations and the Grading Gap

Equitable Grading Practices (EGP) prioritize accuracy, bias resistance, and motivation by focusing on student proficiency rather than behavior or learning speed [1]. Unlike traditional 0–100 scales, which are mathematically biased toward failure and highly sensitive to outliers like "toxic" zeros [6, 7], EGPs employ mastery-based models such as specifications grading [8] or

[1] Instructure, Inc. *Canvas LMS*. `https://www.instructure.com/canvas`
[2] Moodle HQ. *Moodle*. `https://moodle.org`

five-point scales [6]. A fundamental tenet of EGP is the most recent performance logic, where subsequent demonstrations of competency supercede earlier attempts [1]. However, because programming concepts are tightly integrated into knowledge dependency graphs, failure to master foundational topics can lead to a permanent "LMS Bottleneck" where stagnant mean scores fail to reflect a student's eventual proficiency [9, 10, 11].

3.2 Technical Infrastructure and the Implementation Gap

The primary barrier to EGP adoption is that the majority of Learning Management Systems (LMS) are architecturally designed for point-based accumulation [12, 13]. This technological momentum forces instructors into linear, time-dependent grading models that conflict with iterative mastery. Current solutions are often fragmented; while standalone applications like the Mastery Learning App or interactive textbooks like OpenDSA successfully implement EGPs, they often lack seamless LMS integration, requiring manual grade exports or supporting only raw score pass-throughs [14, 15]. Furthermore, ad hoc tools developed by teaching assistants frequently lack documentation, leading to a cycle of "reinventing the wheel" across institutions rather than building adaptable, shared solutions [16, 17].

3.3 Protocol Evolution: Bridging LTI 1.1 and 1.3

Learning Tools Interoperability (LTI) provides the standard for connecting external tools to an LMS securely [18]. While the legacy LTI 1.1 protocol is restricted to passing single numeric grades [19], the modern LTI 1.3 Advantage standard introduces Assignment and Grade Services (AGS), allowing for sub-grades, categorical tags, and detailed feedback [18]. Despite these advancements, the infrastructure remains underutilized; prominent tools like Gradescope often transmit only raw scores to the LMS, leaving rich feedback data siloed within the external application [20]. This disconnect suggests that a protocol alone is insufficient; there is a critical need for an architectural layer, such as the proposed broker, that can translate mastery-based data into LMS-compatible formats using a Zero-Trust, full-stack framework.

4 Methods

The methodology for the EGP Broker focuses on the development of a middleware architecture designed to decouple grading logic from the technical constraints of modern Learning Management Systems (LMS). By acting as a logic gateway, the system enables the implementation of Equitable Grading Practices (EGP) without requiring modifications to the LMS core or the external assessment tools.

4.1 System Architecture and Rationale

The proposed system employs a "middleware" proxy designed to bridge the functional disparity between an institutional LMS, such as Canvas or Moodle, and external tools like OpenDSA[3] and Gradescope[4]. The fundamental motivation is to address the LMS Bottleneck, where standard

[3]Virginia Tech. *OpenDSA e-Textbook Project*. `https://opendsa.org/`

[4]Turnitin, LLC. *Gradescope*. `https://www.gradescope.com/`

platforms default to point-based accumulation. This architecture enables educators to use alternative models, such as mastery-based or categorical assessments, by providing a flexible interface for configuring grading scales independently of the LMS gradebook.

The system is implemented as a full-stack web application using the Nuxt 4[5] framework and the Nitro[6] server engine. PostgreSQL[7] serves as the central repository for platform registrations, user identity mappings, and specific grading rules. This stack was selected to execute cryptographic operations and database look-ups asynchronously, ensuring that the proxy intervention does not introduce perceptible latency for the student.

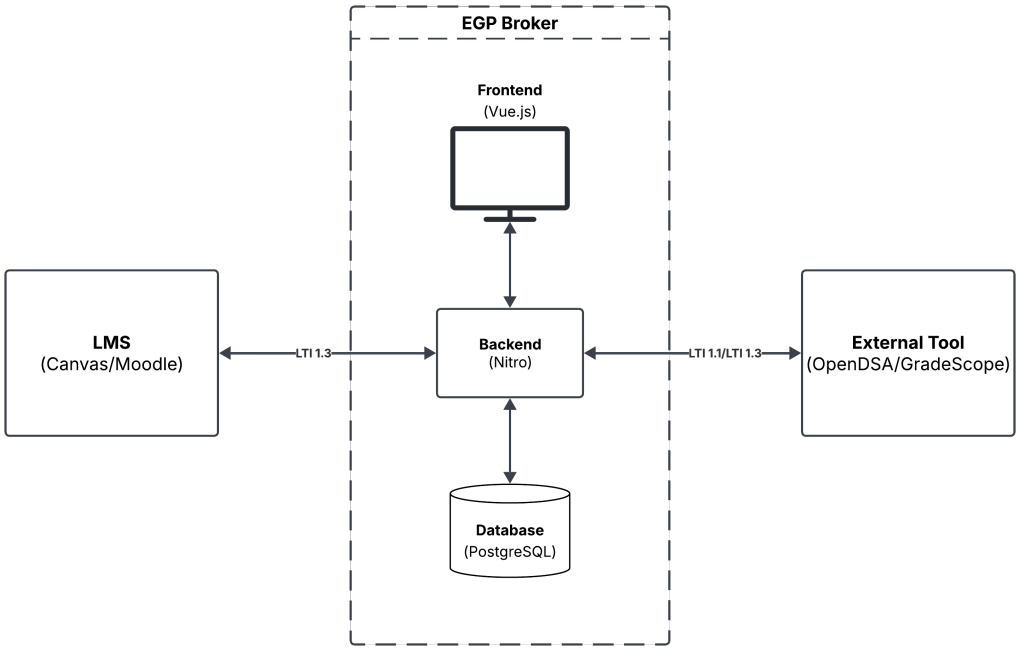

Figure 1: System architecture of the EGP Broker, illustrating its position as a logic gateway between the LMS and external tools.

4.2 Technical Implementation and Protocol Bridging

Communication across all components is governed by the Learning Tools Interoperability (LTI) standard. While built around the modern LTI 1.3/Advantage standard, utilizing OpenID Connect (OIDC) and Assignment and Grade Services (AGS), the broker is designed with protocol-agnostic compatibility. It serves as a bridge for legacy infrastructure, capable of intercepting an LTI 1.3 launch from a modern LMS and adjusting it for tools still utilizing LTI 1.1.

The proxy's operational flow is defined by three primary phases. During the LTI Launch Brokering phase, the broker validates the incoming request from the LMS and performs Dynamic Payload Modification, replacing the standard LMS grade-return URL with the broker's own

[5]Nuxt Team. *Nuxt*. Version 4.0. `https://nuxt.com/`
[6]Unjs. *Nitro Server Engine*. Version 2.10. `https://nitro.unjs.io/`
[7]PostgreSQL Global Development Group. *PostgreSQL*. Version 16.2. `https://www.postgresql.org/`

interception API (`/proxy/grade-passback`) to ensure all performance data flows through the proxy. The Identity Mapping and Security phase involves a session mapping layer through an `LtiIdentity` table that links external tool IDs back to the original LMS security context. Within this phase, a Zero-Trust security model is enforced for all grade submissions, mandating cryptographic verification via JWKS endpoints or shared secrets rather than relying on source IP addresses. Finally, in the Grade Interception and Transformation phase, the broker intercepts student performance data submitted by a tool and applies instructor-defined EGP rules. This process allows for the conversion of raw data into categorical mastery scores, such as Mastery or Developing, or most-recent-performance updates before the final grade is securely pushed to the LMS AGS endpoint.

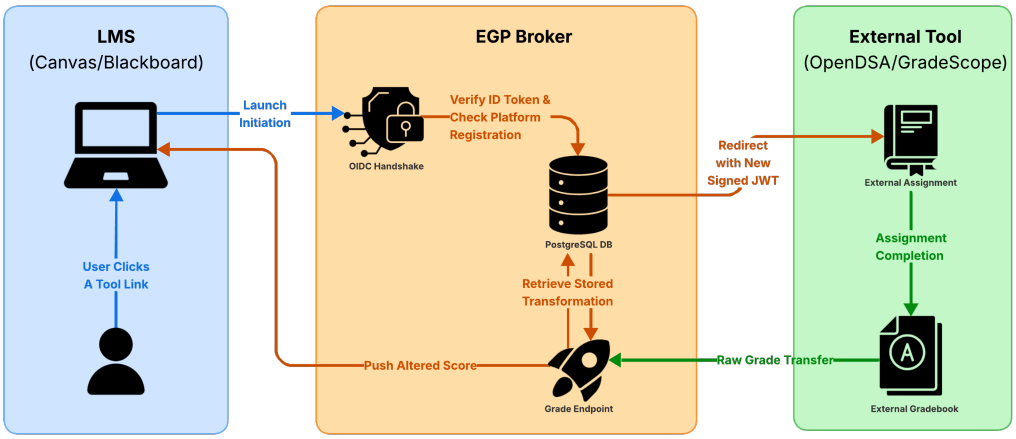

Figure 2: Sequence diagram illustrating the launch brokering and grade transformation process.

4.3 Testing and Validation Procedure

Validation of the proxy was conducted in a controlled environment using a non-production Canvas instance and CodeWorkout as the external assessment tool. The proxy was developed locally and exposed via a secure tunnel (ngrok[8]) to allow the Canvas instance to connect over a public-facing URL.

The validation procedure followed a three-step protocol to confirm end-to-end functionality:

1. **Launch Verification:** Confirming a successful LTI 1.3 launch from Canvas, routed through the proxy, resulting in the correct resource display in the external tool.
2. **Grade Interception:** Confirming that the external tool successfully submits a grade payload to the proxy endpoint and that the proxy correctly retrieves the translation rules from the PostgreSQL database.
3. **Grade Passback Confirmation:** Confirming that the proxy's translation logic executes accurately (e.g., applying a floor grade or categorical mapping) and the final score is

[8]Ngrok, Inc. *ngrok*. `https://ngrok.com/`

recorded correctly in the Canvas Gradebook.

5 Results

This section presents the functional validation of the EGP Broker within a controlled development environment. By evaluating a single student's workflow from a Learning Management System (LMS) launch to the final grade report, we demonstrate the system's ability to bridge disparate protocols and apply pedagogical transformations.

5.1 Experimental Setup: Canvas and CodeWorkout

Validation was conducted using Canvas as the institutional LMS and CodeWorkout as the external assessment tool. CodeWorkout was selected due to its use of the legacy LTI 1.1 protocol and its reliance on raw, points-based scoring. This environment allowed us to test the broker's ability to process an LTI 1.3 launch from Canvas and modify the payload for a legacy consumer.

5.2 LTI Launch and Payload Intercept

The first stage of validation confirms that the broker can transparently intercept and redirect an LTI launch request. During the launch, the broker performs Dynamic Payload Modification by replacing the standard LTI 1.3 Assignment and Grade Services (AGS) lineitem with its own interception endpoint. As shown in the trace below, the broker accurately remapped the `lis_outcome_service_url` to point to the internal proxy URL. Furthermore, the broker re-signed the launch payload using the HMAC-SHA1 signature method required for LTI 1.1 compatibility, utilizing the specific consumer keys stored in the PostgreSQL database.

Listing 1: LTI 1.1 Brokered Launch Log

```
egp-broker-dev | --- OIDC INITIATION ---
egp-broker-dev | Final check - Role: TEACHER Needs Config: true
egp-broker-dev | --- GENERATING REDIRECT ---
egp-broker-dev | ltiParams: { lti_message_type: 'basic-lti-launch-
   request',
egp-broker-dev | lti_version: 'LTI-1p0',...
egp-broker-dev | lis_outcome_service_url:
egp-broker-dev | 'https://egpbroker.share.zrok.io/api/proxy/grade-
   passback',
egp-broker-dev | oauth_consumer_key: '[REDACTED_OAUTH_KEY]',
egp-broker-dev | oauth_signature_method: 'HMAC-SHA1',
egp-broker-dev | oauth_timestamp: '[TIMESTAMP]',
egp-broker-dev | oauth_version: '1.0' }
egp-broker-dev | SIGNATURE: [REDACTED_HMAC_SIGNATURE]
```

5.3 Grade Interception and Transformation Logic

Once the activity is completed, CodeWorkout calculates a raw performance score (e.g., 8/10 tests passed). Instead of transmitting this 80% score directly to the LMS, the tool sends an XML-based

`replaceResult` request to the broker's interception API. The broker then executes the following sequence:

- **Context Retrieval:** The broker uses the `lis_result_sourcedid` from the incoming XML to resolve the original LTI 1.3 session and identity mappings stored in the `LtiIdentity` table.
- **EGP Transformation:** The proxy retrieves the instructor's specific grading schema—such as a Mastery/Non-Mastery binary or a categorical rubric—from the `Assignment` table and applies the logic to the raw score.
- **LTI Advantage Uplift:** The broker uplifts the transformed data into a modern LTI 1.3 `Score` object and transmits it to the Canvas AGS endpoint using an OAuth 2.0 Bearer token.

5.4 LMS Gradebook Integration

The success of the EGP Broker is confirmed when the transformed grade appears within the native Canvas environment. The gradebook reflects the value defined by the instructor's transformation rules rather than the raw percentage from CodeWorkout. This validation confirms that the broker effectively decouples assessment scoring from grade calculation, providing a seamless experience that is entirely invisible to the student.

6 Discussion

The functional validation of the EGP Broker demonstrates that a middleware proxy is a robust technical solution to the systemic limitations that currently preclude the adoption of Equitable Grading Practices (EGP). By serving as an abstraction layer, the broker provides several significant implications for both pedagogy and technical infrastructure.

6.1 Decoupling Instructional Intent from Infrastructure

The primary contribution of this work is the decoupling of instructional goals from the rigid data structures of the Learning Management System (LMS). Most LMS platforms are architected around a linear, points-based grading model that reduces comprehensive categorical feedback into a simple numerical percentage. The EGP Broker preserves the pedagogical integrity of EGPs while performing the necessary translation into the numerical values required by institutional gradebooks for final reporting.

Furthermore, this middleware approach significantly reduces the administrative burden associated with alternative grading. Without such a tool, instructors in high-enrollment computer science courses must often manually export raw scores, translate them in spreadsheets, and re-enter them into the LMS by hand. The EGP Broker automates this workflow, transforming equitable grading from a high-effort manual task into a sustainable, scalable practice.

6.2 Security and Architectural Integrity

By implementing a brokered token exchange, the system extends rather than weakens the security of the LTI 1.3 standard. The Zero-Trust verification of external tool signatures ensures that the proxy upholds production-grade security standards.

From an interoperability standpoint, the architecture supports a one-to-many mapping. A single instance of the EGP Broker can manage a wide range of tools like OpenDSA and CodeWorkout, providing a centralized point of configuration for grading logic across an entire curriculum.

7 Limitations

While the system successfully manages the LTI lifecycle, two primary limitations were identified. First, the current implementation facilitates a one-to-one relationship between an external activity and an LMS gradebook column. Many alternative models, such as Specifications Grading, require more sophisticated accumulative logic where a final module grade is determined by a combination of diverse signals (e.g., combining scores from both OpenDSA and Gradescope). Under the current architecture, these signals remain siloed.

Second, the current system focuses primarily on score transformation and does not yet fully integrate structured feedback. While the architecture supports these extensions, a comprehensive model for aligning qualitative feedback with specific grading categories remains a subject for future development.

8 Conclusion and Future Work

This research developed and validated a middleware broker that enables EGP use through the LTI 1.3 protocol. By employing the Nuxt 4 framework and Nitro server engine, the system provides a high-performance proxy capable of both modern LTI 1.3 and legacy LTI 1.1 brokering. The integration of a PostgreSQL database allows for the persistent storage of complex grading rules, enabling the dynamic application of floor grading and categorical mastery mapping without requiring modifications to the LMS or the assessment tools.

Future work will focus on broadening the pedagogical scope of the broker. A key priority is the implementation of a user-friendly, no-code interface in Vue.js to allow non-technical instructors to define their own transformation rules. Technically, we plan to incorporate support for the SPLICE standards to better align with computer science education research and data standards. Additionally, we aim to implement advanced aggregation logic to support multi-tool grading models and conduct large-scale user studies to evaluate the impact on instructor workload and student perceptions of fairness. Ultimately, the EGP Broker provides a scalable framework that enables educators to prioritize demonstrable learning mastery over the raw accumulation of points.

References

[1] J. Feldman, *Grading for Equity: What It Is, Why It Matters, and How It Can Transform Schools and Classrooms*, 2nd ed. Corwin Press, 2024.

[2] L. B. Nilson, *Specifications Grading: Restoring Rigor, Motivating Students, and Saving Faculty Time*. Stylus Publishing, 2015.

[3] A. Decker, S. H. Edwards, B. Edmison, M. A. Pérez-Quiñones, and A. Rorrer, "A call for critical technology to enable innovative and equitable grading practices," in *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE TS 2026)*. New York, NY, USA: ACM, 2026.

[4] A. Decker, S. H. Edwards, B. Edmison, M. A. Pérez Quiñones, and A. Rorrer, "A call for critical technology to enable innovative and alternative grading practices," in *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V.1*, ser. SIGCSE TS 2026. New York, NY, USA: Association for Computing Machinery, 2026, p. 260–266. [Online]. Available: https://doi.org/10.1145/3770762.3772530

[5] IMS Global Learning Consortium, "Lti 1.3 assignment and grade services specification," n.d., accessed through the 1EdTech website. [Online]. Available: https://www.imsglobal.org/spec/lti-ags/v2p0/

[6] D. Reeves, "The case against the zero," *Phi Delta Kappan*, vol. 86, pp. 324–325, 12 2004.

[7] T. R. Guskey, "Bound by tradition: Teachers' views of crucial grading and reporting issues," in *Annual Meeting of the American Educational Research Association*, San Francisco, CA, United States, April 2009.

[8] W. J. Rapaport, "A new approach to grading: Triage grading," in *Proceedings of the 54th Annual Southeast Regional Conference*. ACM, 2016, pp. 53–58.

[9] A. Robins, "Learning edge momentum: a new account of outcomes in CS1," *Computer Science Education*, vol. 20, no. 1, pp. 37–71, 2010.

[10] C. Szabo, M. C. Parker, M. Friend, J. Jeuring, T. Kohn, L. Malmi, and J. Sheard, "Models of mastery learning for computing education," in *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE TS 2025. New York, NY, USA: Association for Computing Machinery, 2025. [Online]. Available: https://doi.org/10.1145/3641554.3701868

[11] T. R. Guskey, "The case against percentage grades," *Educational Leadership*, vol. 71, no. 1, pp. 68–72, 2013.

[12] Instructure, *Canvas Basics Guide*, 2026, accessed: 2026-04-16. [Online]. Available: https://community.canvaslms.com/t5/Canvas-Basics-Guide/tkb-p/basics

[13] Anthology Inc., *Blackboard Learn Help for Instructors*, 2026, accessed: 2026-04-16. [Online]. Available: https://help.blackboard.com/Learn/Instructor

[14] E. Tuson and T. Hickey, "Mastery learning with specs grading for programming courses," in *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2023. New York, NY, USA: Association for Computing Machinery, 2023, p. 1049–1054. [Online]. Available: https://doi.org/10.1145/3545945.3569853

[15] C. A. Shaffer, V. Karavirta, A. Korhonen, and T. L. Naps, "Opendsa: beginning a community active-ebook project," in *Proceedings of the 11th Koli Calling International Conference on Computing Education Research*, ser. Koli Calling '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 112–117. [Online]. Available: https://doi.org/10.1145/2094131.2094154

[16] A. Brooks-Ramirez and L. Yan, "By and for teaching assistants: Homegrown tools in computing classrooms," in *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V.2*, ser. SIGCSE TS 2026. New York, NY, USA: Association for Computing Machinery, 2026, p. 1249–1250. [Online]. Available: https://doi.org/10.1145/3770761.3777245

[17] J. Blanchard, J. R. Hott, V. Berry, R. Carroll, B. Edmison, R. Glassey, O. Karnalim, B. Plancher, and S. Russell, "Stop reinventing the wheel! promoting community software in computing education," in *Proceedings of the 2022 Working Group Reports on Innovation and Technology in Computer Science Education*, ser. ITiCSE-WGR '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 261–292. [Online]. Available: https://doi.org/10.1145/3571785.3574129

[18] E. Consortium, "Learning tools interoperability core specification v1.3," 2023, accessed: 2026-04-17. [Online]. Available: https://www.imsglobal.org/spec/lti/v1p3

[19] ——, "Learning tools interoperability (lti)® basic outcomes v1.1," 2012, accessed: 2026-04-17. [Online]. Available: https://www.imsglobal.org/spec/lti-bo/v1p1

[20] A. Singh, S. Karayev, K. Gutowski, and P. Abbeel, "Gradescope: A fast, flexible, and fair system for scalable assessment of handwritten work," in *Proceedings of the Fourth (2017) ACM Conference on Learning @ Scale*, ser. L@S '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 81–88. [Online]. Available: https://doi.org/10.1145/3051457.3051466