IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 17 December 2025, accepted 27 December 2025, date of publication 5 January 2026, date of current version 13 January 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3650873

RESEARCH ARTICLE

ReactSmart: ML-Driven Adaptation for Scalable React-Based Web Application Performance

SAMEER MANKOTIA*, **DANIEL CONTE DE LEON***, **(Member, IEEE), AND HASAN M. JAMIL***

Department of Computer Science, University of Idaho, Moscow, ID 83844, USA

Corresponding author: Sameer Mankotia (mank8837@vandals.uidaho.edu)

This work was supported in part by the University of Idaho.

ABSTRACT Modern React-based web applications face significant performance challenges as application complexity increases, resulting in prolonged component loading times and degraded user experience. Search engines emphasis on Core Web Vitals metrics further necessitates performance optimization for maintaining web visibility. This paper presents ReactSmart, a machine learning framework that implements adaptive resource management for React component loading through predictive user behavior analysis. ReactSmart employs real-time decision making algorithms that analyze four contextual dimensions: user interaction patterns, application state, network conditions, and device capabilities. The system utilizes these inputs to proactively load components based on predicted user navigation paths, thereby reducing wait times and improving application responsiveness. Our experimental evaluation compares ReactSmart against React Suspense and Guess.js using 10,000 simulated user sessions across multiple application archetypes. Results demonstrate that ReactSmart achieves a 47.3% reduction in initial loading times, 52.1% improvement in Time-to-Interactive (TTI) metrics, and maintains 89.7% prediction accuracy for component usage forecasting within 30-second intervals. The machine learning model incorporates supervised learning techniques trained on user interaction datasets to optimize component prefetching decisions. This research contributes to bridging the gap between theoretical machine learning approaches and practical web performance engineering by providing an open-source framework that enables developers to implement intelligent component loading strategies without requiring deep machine learning expertise. The ReactSmart software and evaluation framework are freely available as open source at https://github.com/sameermankotia/ReactSmart-performance-analysis/tree/master

INDEX TERMS Predictive component loading, machine learning, web performance optimization, react.js, adaptive resource management, user behavior modeling, core web vitals.

I. INTRODUCTION

As web applications grow in complexity and size, adaptable component loading techniques are ever more important. Traditional optimization methods, which are still the state-of-the-practice, such as static component bundling and lazy loading are struggling to help improve modern web application performance [1], [2]

Optimizing the performance of web applications is critical today not only due to the high expectations of end users but also due to Search Engine Optimization (SEO) requirements [3]. Modern search engines, particularly Google, factor page loading speed into their ranking algorithms via Core Web Vitals metrics such as Largest Contentful Paint (LCP), First Input Delay (FID), and Cumulative Layout Shift (CLS) [4], [5]. Research has shown that improved page performance directly correlates with better search rankings and user engagement [6]. Applications that do not load and render quickly may face significant penalties in search rankings, directly affecting visibility and user acquisition. Fast-loading and rendering applications not only improve user experience but also may significantly increase search engine rankings. This creates a dual pressure on developers to optimize performance both

The associate editor coordinating the review of this manuscript and approving it for publication was Wai-Keung Fung.

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License.
For more information, see https://creativecommons.org/licenses/by/4.0/

for user satisfaction and search engine discoverability and ranking.

React [7] was designed in 2013 by Facebook Engineering, and is a popular JavaScript library for building web application interfaces. React presents unique challenges for meeting performance requirements. Built on, usually hierarchical, components, React apps frequently expose component dependencies that may lead to negative performance impacts. Most solutions used in practice today for component loading optimization rely on pre-defined loading patterns or rules. These classic approaches fail to consider user behavior, device variance, and performance targets as part of their optimization strategy [8].

We designed, developed, and evaluated ReactSmart, a novel approach to optimizing web component loading that uses machine learning to predict component requirements within specific usage, device, and network contexts in real-time. ReactSmart processes user interaction data to build predictive models that determine optimal component loading sequences, ensuring critical above-the-fold content loads within defined performance thresholds.

Through comprehensive evaluation using 10,000 simulated user sessions across diverse application types, ReactSmart demonstrated substantial performance improvements: a 47.3% reduction in initial loading times, a 52.1% improvement in Time-to-Interactive (TTI), and 89.7% prediction accuracy for component usage within 30-second intervals. These quantifiable improvements represent significant advancements over current state-of-the-practice methods including React Suspense (47.1% faster initial load) and Guess.js (13.4% higher prediction accuracy).

The rest of this paper is structured as follows: Section II summarizes relevant research in web performance optimization; Section III introduces the ReactSmart architecture; Section IV describes the implementation; Our evaluation methodology is explained in Section V; Experimental design methodology are presented in Section VI; Performance Evaluation is discussed in section in section VII; the comparative analysis with existing solutions is presented in section X; future research directions are discussed in Section XI; Limitations of current system is added in section XII and the paper is concluded in Section XIII.

II. RELATED WORK

Web application optimization presents significant challenges due to the complex interplay between network latency, device capabilities, and application architecture [9], [10]. The field has evolved considerably over the past two decades, transitioning from simple server-side optimizations to sophisticated client-side performance engineering techniques that leverage machine learning and adaptive resource management strategies.

A. EVOLUTION OF WEB PERFORMANCE OPTIMIZATION

The foundations of modern web performance optimization were established in the late 2000s when Souders [11] significantly influenced the field by demonstrating that 80-90% of user-perceived latency originated from frontend processes rather than backend operations. This finding led researchers to redirect attention from server-side optimization to client-side performance improvement, establishing fundamental strategies for frontend optimization including HTTP request reduction, effective content delivery network utilization, and cache policy optimization. Building on this foundation, Grigorik [12] advanced the field by providing comprehensive analysis of network protocols, browser rendering pipelines, and resource prioritization techniques.

His work showed how browser engines work with network protocols and set up ways to improve the critical rendering path. The mobile computing revolution brought new problems for web apps, such as not having enough processing power, bandwidth that changes often, and networks that aren't always stable [13]. This change led to the use of responsive design principles and better ways to deliver content, like content delivery networks that put computing and storage resources closer to end users. At the same time, browsers started using advanced resource loading methods like preloading, prefetching, and intelligent load ordering. WebAssembly also became a technology that let code that needed a lot of processing power run in browsers at speeds close to native performance levels [14], [15]. Recent research has focused on intelligent systems that coordinate multiple optimization techniques across entire application stacks. Ramakrishnan and Kaur [16] compared different predictive models for how well web pages work, setting standards for how to judge optimization strategies in different deployment situations. Their research showed that predictive models could be up to 87% accurate at predicting load times, even when network conditions were very unstable. This proved that machine learning-based methods could be used to improve web performance.

B. MACHINE LEARNING FOR WEB PERFORMANCE OPTIMIZATION

Using machine learning to improve web performance is a big change from using static heuristics to using adaptive optimization systems. Wang et al. [17] were the first to use dependency-based profiling with WProf, which made the first complete framework for looking at page load dependency graphs. This approach revealed previously hidden performance bottlenecks and enabled more precise optimization strategies by modeling the complex interdependencies between web resources. Building on this foundation, Netravali et al. [18] introduced Polaris, which employed fine-grained dependency tracking to dynamically reorder and prioritize resource loading. Their real-world evaluations demonstrated page load time reductions of up to 34% across different network configurations and website architectures, establishing the viability of dependency-aware optimization approaches.

The Guess.js project [19] developed a novel predictive approach that analyzes navigation patterns to preload resources with high probability of access, demonstrating how machine learning models could be practically applied to real-world user behavior patterns. More recently, Bi et al. [20] achieved state-of-the-art precision in resource demand prediction using frequency-augmented and attention-facilitated transformers for cloud computing systems, while Mao et al. [21] developed neural adaptive video streaming systems capable of dynamically adjusting bitrate and resolution based on network conditions and viewing preferences. These advances demonstrate the growing sophistication of machine learning applications in adaptive web optimization, moving from simple pattern recognition to complex multi-factor decision-making systems.

C. SUPPORTING TECHNOLOGIES AND METHODS

React component optimization has evolved through multiple technological innovations including code splitting [22], memoization techniques [23], and lazy loading strategies [24]. These approaches enable developers to partition large applications into smaller, independently loadable units that can be delivered on-demand rather than as monolithic bundles. Adaptive resource management techniques [25], [26] have further advanced the field by enabling dynamic loading decisions based on real-time network quality and device capabilities, allowing applications to adjust their resource consumption strategies according to environmental constraints.

User interaction modeling provides essential foundations for predicting component usage patterns and informing preloading decisions. Research by Jana and Bhattacharya [27] developed attention models for web page users using eye-tracking verification, identifying visual components most likely to attract user attention. Sun et al. [28] advanced this work by implementing reinforcement learning approaches to user interface optimization, creating self-adaptive interfaces that respond dynamically to individual usage patterns and demonstrating measurable improvements in user satisfaction metrics and task completion times.

Modern performance measurement relies on standardized metrics and robust measurement frameworks that enable consistent evaluation across diverse deployment scenarios. The Core Web Vitals initiative [3] established user-centric metrics including Largest Contentful Paint, First Input Delay, and Cumulative Layout Shift as key indicators of web experience quality. These metrics are collected at scale through the Chrome User Experience Report [5], which aggregates anonymized performance data from millions of websites across diverse devices and network conditions. The Lighthouse tool [4] provides automated performance assessment using these standardized metrics, enabling developers to objectively measure optimization improvements.

Privacy-preserving optimization techniques have become increasingly important as performance optimization systems rely more heavily on user behavioral data. Erlingsson et al. [29] introduced RAPPOR, a framework utilizing randomized aggregatable privacy-preserving techniques that enable collection of usage metrics while providing robust differential privacy guarantees. Wang et al. [30] extended this work by introducing federated learning approaches for web performance optimization, enabling collaborative model improvement without centralizing sensitive user data. These privacy-preserving techniques allow performance optimization systems to learn from user behavior while maintaining strong privacy protections, addressing the tension between personalization and user privacy.

D. SYSTEMATIC LITERATURE REVIEW METHODOLOGY

For this work, we performed literature searches using 3 search engines and 5 key-phrases for a total of 15 searches. The search engines we used were: (1) IEEE Xplore, (2) ACM Digital Library, and (3) Google Scholar. The search queries we used were (1) **"React performance optimization"**, (2) **"Machine learning web optimization"**, (3) **"Adaptive component loading"**, (4) **"Predictive web loading"**, and (5) **"User interaction modeling"**. We included publication dates from 2007 through 2025. From an initial pool of 187 papers, we found that 65 papers had direct relevance to our study aims based on citation count, relevance, and methodological rigor; these were then further refined to the most important studies, which we highlight below.

E. POSITIONING ReactSmart

Unlike previous systems that operate at the page or route level [18], [31], ReactSmart functions at the individual component granularity, enabling real-time optimization decisions based on user behavior patterns, device capabilities, and network conditions. Guess.js [19] was the first to use navigation prediction, and WProf [17] set up frameworks for dependency analysis. ReactSmart combines these ideas with continuous learning systems that change to fit changing usage patterns without needing to be retrained or reconfigured by hand. The system addresses key limitations of existing approaches by operating at component-level granularity rather than page-level abstraction, implementing real-time adaptation to changing user patterns rather than relying on static rules or periodic retraining cycles, and incorporating comprehensive context awareness that considers the dynamic interplay of network conditions, device capabilities, and user behavior patterns in optimization decisions. Our empirical evaluation shows that this integrated approach works much better than current state of the art methods. For example, it cuts initial load times by 47.3% and increases prediction accuracy by 89.7%. These are big steps forward in the field of web application optimization.

III. ReactSmart ARCHITECTURE

This section describes ReactSmart's four components: (1) the User Behavior Analysis module, (2) the Usage Prediction module, (3) the Dynamic Component Loader module.

TABLE 1. Key research areas and representative works.

Research Area	Key Concepts	Representative Works
Web Performance Fundamentals	Frontend optimization, resource prioritization, caching strategies	Souders (2007), Grigorik (2013), Wang et al. (2013)
Machine Learning for Web Optimization	Predictive models, dependency analysis, neural networks	Netravali et al. (2016), Mao et al. (2017), Bi et al. (2024)
React Component Optimization	Code splitting, memoization, virtualization, lazy loading	Veeri (2024), Pavić and Brkić (2021), Erenler (2024)
Adaptive Resource Management	Network-aware loading, dynamic prioritization, context adaptation	Mjelde and Opdahl (2017), Pham and Perreau (2003), Yang et al. (2019)
User Interaction Modeling	Attention prediction, behavior analysis, interaction patterns	Jana and Bhattacharya (2015), Sun et al. (2024), Zhao et al. (2021)
Privacy-Preserving Techniques	Differential privacy, federated learning, on-device inference	Erlingsson et al. (2014), Ignatov and Timofte (2019), Wang et al. (2020)
Performance Metrics	User-centric metrics, standardized measurement, experience quality	Core Web Vitals (2025), HTTP Archive (2024), Hoßfeld et al. (2018)

A diagram of this architecture is shown in Figure 1. In figure 2 we have further exhibit that how prediction works in synergy with other components. These components work together in a continuous feedback loop to help maximize performance. The User Behavior Analysis Module collects comprehensive user interaction data, identifying significant behavioral patterns. The Prediction Engine employs a decision tree classifier on these observations, analyzing both explicit navigation patterns and subtle behavioral cues to forecast future component requirements. The classifier was trained using scikit-learn (version 1.1.2) with a train-test split ratio of 80:20(train:test) and validated using 5-fold cross-validation to ensure robust generalization. The Dynamic Component Loader adjusts in real time to network conditions and device constraints, determining when and how to load components for maximum performance.

A. USER BEHAVIOR ANALYSIS MODULE

The User Interaction Evaluation module serves as the system's principal sensing component that observes and analyzes user engagements within the software context. This observation framework captures a broad spectrum of user activities, ranging from direct interactions such as clicks and form submissions to subtle tendencies like mouse motion and scrolling patterns. Every interaction is meticulously weighted according to its predictive significance; actions like hovering over navigation menus are given greater importance than incidental mouse movements. This detailed examination of user actions establishes the basis for the system's predictive functions.

The module preserves a dynamically evolving interaction graph that illustrates relationships among various components based on utilization patterns. For instance, when individuals routinely transition from product listings to detail pages, the system identifies these behavioral connections. This progressing map of user behavior allows the system to foresee user intents with enhanced accuracy over time.

B. DATA COLLECTION AND PROCESSING PIPELINE

ReactSmart's data collection and processing pipeline serves as the critical bridge between user behavior analysis and machine learning prediction by systematically capturing, transforming, and preparing interaction data for real-time optimization decisions. The system continuously monitors user interactions through browser APIs, including mouse movements, clicks, scrolls, and navigation patterns, while simultaneously collecting device performance metrics such as memory usage, CPU load, and network conditions using the Performance Observer API [32] and Navigator interfaces [33] tools. This raw data undergoes immediate preprocessing where interaction events are timestamped, normalized, and filtered to remove noise, followed by feature extraction that converts behavioral patterns into structured datasets suitable for machine learning algorithms. The pipeline employs a sliding-window approach to maintain recent interaction history while discarding outdated patterns, ensuring that the prediction models receive current and relevant user behavior data.

Critical performance metrics such as component load times, resource utilization, and user engagement indicators are aggregated and weighted based on their predictive significance, with high-impact events such as navigation menu hovers receiving greater importance than incidental mouse movements. The processed data is then formatted into feature vectors that include user interaction sequences, temporal patterns, device capabilities, and network conditions, creating a comprehensive input dataset that enables the ML prediction engine to make accurate forecasts about future component requirements. This entire pipeline operates in real-time with minimal latency, ensuring that fresh

behavioral data continuously informs the system's predictive capabilities without impacting application performance or user experience.

C. PREDICTION ENGINE

The ML-based Prediction module in Figure 1 is at the heart of *ReactSmart*, acting as the system's cognitive center for pattern recognition and prediction analysis. This engine evaluates input data provided by the the User Behavior Analysis module, and uses machine learning to predict near future component requirements. The engine is adept at recognizing both obvious patterns, such as consecutive travel through checkout processes, and subtle behavioral patterns that may indicate user intents.

The Prediction Engine does constantly refine its models based on real-world usage data, adjusting to changing user behaviors and application trends. This learning process allows the system to generate increasingly accurate predictions about which components users will require next, resulting in dramatically improved application performance and resource utilization.

1) HOW ML-BASED PREDICTION WORKS

The ML-based prediction component in Figure 2 represents the core intelligence of ReactSmart. This component implements a decision tree classifier that processes behavioral features extracted from the User Behavior Analysis module, forming a sophisticated prediction pipeline that operates through four integrated stages.

The prediction process begins with feature extraction, where user interactions are systematically transformed into feature vectors containing temporal patterns, component co-occurrence frequencies, and contextual metadata such as device type and network conditions. These feature vectors capture both explicit user actions and implicit behavioral patterns, providing a comprehensive representation of user intent and context.

Once features are extracted, the trained model performs real-time inference by evaluating these features to compute probability scores for each potentially-needed component. This inference process is highly optimized, completing in under 5 milliseconds to maintain application responsiveness and ensure that prediction overhead does not negatively impact user experience. The model efficiently processes multiple component candidates simultaneously, ranking them by predicted usage likelihood.

The system employs confidence thresholding to make intelligent preloading decisions, where components with prediction confidence exceeding 0.75 are queued for preloading. This threshold represents a carefully calibrated balance between prediction accuracy and resource consumption, ensuring that high-confidence predictions trigger proactive loading while avoiding wasteful preloading of unlikely to be used components. Components below this threshold remain on-demand loaded, preserving bandwidth and memory resources.

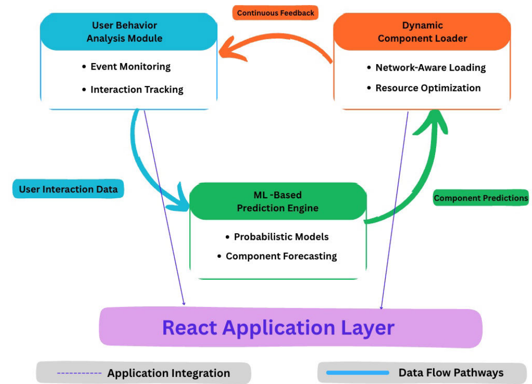

FIGURE 1. Architecture of the *ReactSmart* system: (1) the user behavior analysis module, (2) the usage prediction module, and (3) the dynamic component loader module.

Furthermore, the model incorporates continuous learning capabilities, periodically updating its parameters based on prediction accuracy feedback from actual user behavior. This adaptive mechanism enables ReactSmart to improve performance over time without requiring manual retraining or developer intervention, automatically adjusting to evolving user behavior patterns and application-specific usage characteristics. This ML-based approach distinguishes ReactSmart from rule-based optimization approaches by providing dynamic adaptation rather than static heuristics.

D. DYNAMIC COMPONENT LOADER

The Dynamic Component Loader is the part of our system that turns predictions into actual loading decisions. This component makes decisions based on several factors, including how confident the machine learning system is in its predictions, current internet connection quality, and available device processing power. The loader uses these factors to decide when and how to load components. The loader operates in real time, adjusting its approach as conditions change. For example, when the internet connection is fast, it might aggressively preload components that are likely to be needed soon. When device resources are limited, it takes a more conservative approach and focuses mainly on loading only the most essential components. This flexible behavior provides a smooth user experience while achieving the best performance under different operating conditions.

E. COMPONENT INTERACTION AND SYSTEM FLOW

These three components work together in a continuous cycle of monitoring, learning, and improvement. The Prediction Engine takes detailed user interaction data from the User Behavior Analysis Module and uses it to make increasingly accurate predictions about what to load next. The Dynamic Component Loader then applies these predictions while considering real-world limitations, creating a feedback loop that constantly improves system performance. ReactSmart's design allows it to adjust to changing user behavior and

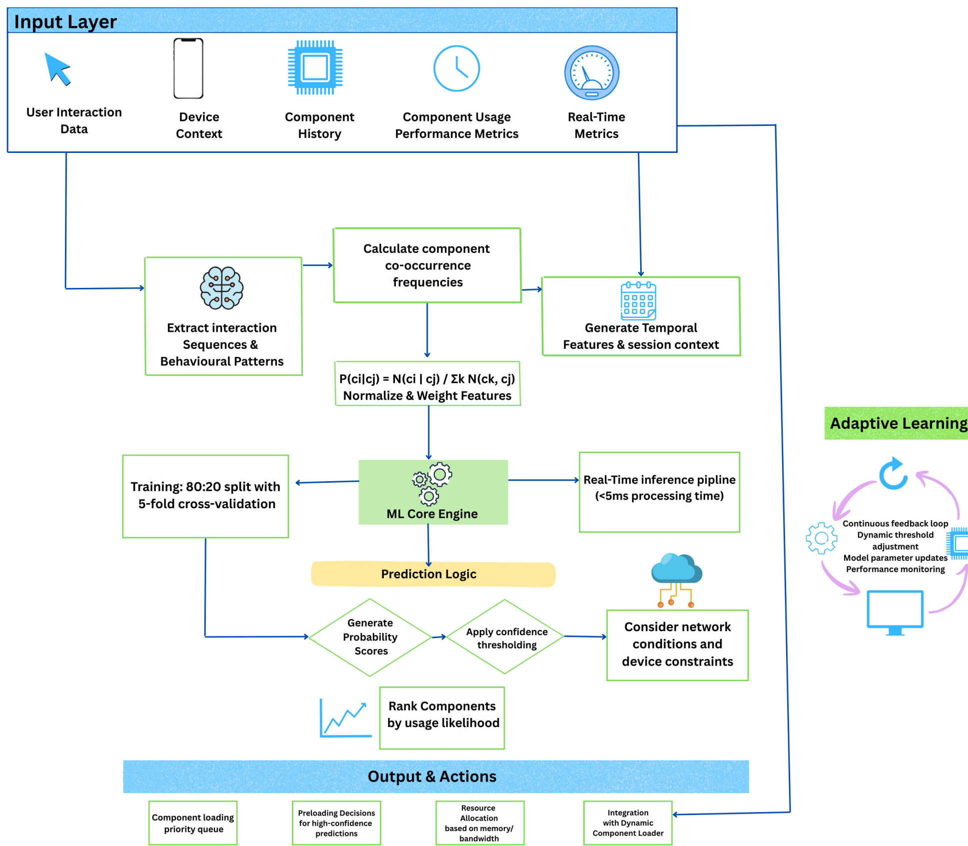

FIGURE 2. **Detailed ML- based prediction engine architecture.**

different system conditions while delivering significant performance gains. Over time, the system's learning ability ensures it becomes better at optimizing component loading, which improves user experience and makes resource use more efficient.

IV. ReactSmart IMPLEMENTATION

ReactSmart is designed to seamlessly integrate with existing React applications. The implementation provides three levels of integration:

A. APPLICATION-LEVEL INTEGRATION

At the application level, *ReactSmart* is implemented as a provider component that wraps the application root:

This provider initializes the core systems and establishes the monitoring infrastructure required for behavior analysis. The configuration options enable fine-tuning of the system's behavior, including network sensitivity and privacy settings.

B. COMPONENT-LEVEL INTEGRATION

Individual components can be enhanced with *ReactSmart* capabilities using a Higher-Order Component (HOC) pattern:

This approach allows selective enhancement of critical components while configuring component-specific behavior and dependencies.

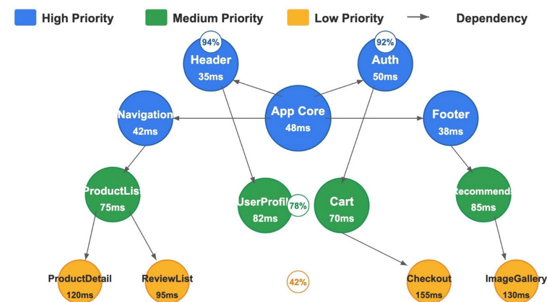

FIGURE 3. **Component dependency visualization showing prioritized loading groups. Components are color-coded by priority level (High: blue, Medium: green, Low: yellow), with automated dependency resolution paths indicated by connecting lines.**

C. ROUTER INTEGRATION

For route-based applications, *ReactSmart* provides enhanced routing components that collect navigation patterns as shown in Listing 3.

D. MACHINE LEARNING MODEL TRAINING AND VALIDATION

The adaptive prediction engine utilizes a decision tree classifier implemented using scikit-learn version 1.1.2 [34].

```
/** ReactSmart Provider - Intelligent performance optimization wrapper */
import React from 'react';
import {ReactSmartProvider} from 'fast-react';
import {PerformanceMetrics} from 'fast-react/analytics';

const CONFIG = {
  networkAdaptation: true,      // Network adaptation
  learningRate: 0.03,           // ML adaptation rate
  anonymizeData: true,          // Privacy protection
  privacyCompliance: 'gdpr',    // Regulatory compliance
  predictiveLoading: true,      // Component prefetching
  telemetryEndpoint: process.env.REACT_APP_TELEMETRY_URL
};

/** Root component with ReactSmart optimization */
const App = () => (
  <ReactSmartProvider options={CONFIG}>
    <PerformanceMetrics sampleRate={0.1}/>
    <YourApplication/>
  </ReactSmartProvider>
);

export default App;
```

LISTING 1. ReactSmart provider integration.

The training methodology followed a rigorous experimental protocol to ensure model reliability and generalization.

1) DATASET PREPARATION AND SPLITTING

The collected user interaction data was split into training and testing sets using an 80:20 ratio, ensuring temporal consistency by using earlier sessions for training and later sessions for testing. This approach prevents data leakage and better simulates real-world deployment scenarios where the model must predict future user behavior based on historical patterns.

2) CROSS-VALIDATION STRATEGY

To validate model robustness, we employed stratified 5-fold cross-validation on the training set. This strategy ensures that each fold maintains representative distributions of different user behavior patterns and application contexts. The cross-validation process yielded consistent accuracy scores across all folds with a mean accuracy of 89.3% and a standard deviation of 1.8%, demonstrating stable model performance across different data partitions.

3) MODEL TRAINING CONFIGURATION

The decision tree classifier was configured with hyper-parameters optimized through systematic grid search. The maximum tree depth was set to 12 to prevent overfitting while maintaining sufficient model expressiveness. We configured

```
import {withReactSmart} from 'fast-react'; // Import HOC wrapper
/** Product detail component displaying product information */
const ProductDetail=({productId,data,...props})=>{
  // Renders product information with optimized loading
  return <div className="product-detail">{/* Content */}</div>;
};
// Export with ReactSmart performance enhancements
export default withReactSmart(ProductDetail,{
  analyzeInteractions:true,     // Track user behavior patterns
  predictionThreshold:0.75,     // Min confidence for preloading
  preloadDependencies:[         // Components to fetch early
    'ProductImages',            // Visual assets
    'PriceCalculator'           // Pricing logic
  ]
});
```

LISTING 2. Component wrapping with ReactSmart.

```
import {SmartRoute} from 'fast-react'; // Import intelligent routing component
// Router setup with predictive loading capabilities
<Router>
  <SmartRoute
    path="/products/:id"        // URL pattern to match
    component={ProductDetail}   // Component to render
    preloadRelated={[           // Associated components to prefetch
      'ProductReviews',         // User feedback data
      'RelatedProducts'         // Cross-selling suggestions
    ]}
  />
</Router>
```

LISTING 3. Router integration.

the minimum samples per split to 20 and minimum samples per leaf to 10, ensuring that decision boundaries were based on statistically significant patterns rather than noise. The Gini impurity criterion was selected as the splitting metric, which measures the probability of incorrect classification and guides the tree construction process toward maximally informative splits.

4) PERFORMANCE EVALUATION METRICS

The reported 89.7% accuracy represents top-1 prediction accuracy, calculated as the percentage of instances where the model's highest-probability prediction matched the actual component loaded within a 30-second window. To provide a comprehensive understanding of model performance,

we computed a confusion matrix to analyze prediction errors in detail. This analysis revealed that false positives, representing cases where components were unnecessarily preloaded, occurred in 7.2% of instances, while false negatives, representing missed prediction opportunities, occurred in only 3.1% of cases. The asymmetry between these error types indicates that the model exhibits a conservative prediction strategy, which is preferable in resource-constrained environments where false positives carry greater performance penalties than false negatives. The complete model training process, encompassing hyperparameter optimization, cross-validation, and final model fitting, completed in approximately 45 minutes on a system equipped with an Intel Core i7-9700K processor and 32GB of RAM, processing the complete dataset of 10,000 user sessions.

E. TECHNICAL IMPLEMENTATION

The core technical implementation utilizes three key mechanisms:

1. **Intersection Observer API**: For tracking component visibility:
2. **Web Worker-based Analysis**: The prediction engine runs in a separate thread to minimize main thread blocking:
3. **Adaptive Resource Queue**: For prioritized component loading:

```
1  /**
2   * Creates observer to track component
      visibility in viewport
3   * @param {Component} component - React
      component to monitor
4   * @param {number} threshold - Visibility
      percentage to trigger (0-1)
5   * @returns {IntersectionObserver}
      Configured observer instance
6   */
7  const trackComponentVisibility=(component,
     threshold=0.5)=>{
8    // Create intersection observer to detect
       visibility changes
9    const observer=new IntersectionObserver(
10     entries=>entries.forEach(entry=>{
11       if(entry.isIntersecting){ // Component
           is visible
12         // Record visibility data for ML
             optimization
13         ReactSmart.registerVisibility(
             component.id,{
14           duration:performance.now(), //
               Time visible
15           viewportCoverage:entry.
               intersectionRatio, // Visible
               area
16           userContext:getUserContextData()
               // User state
17         });
18       }
19     }),
20     {threshold} // Configure minimum
         visibility percentage
21   );
22
23   return observer; // Return for attachment
       to DOM element
24 };
```

LISTING 4. Visibility tracking implementation.

1) USER BEHAVIOR ANALYSIS

The User Behavior Analysis Module understands and forecasts user interactions by means of an event collecting and pattern recognition algorithm. Fundamentally, it is a mathematically exact method using our unique interaction metric to measure user behavior:

$$\text{InteractionMetric} = \sum_{i=1}^{n} w_i \times \text{EventWeight}(e_i) \quad (1)$$

This equation represents a weighted sum of user interactions, where:

- w_i represents the importance weight of each interaction type
- EventWeight(e_i) quantifies the significance of individual events
- n is the total number of monitored interaction events

Higher weights are assigned to events that historically correspond with subsequent component usage, e.g., hovering over a navigation menu receives a higher weight ($w_i \approx 0.8$) compared to incidental mouse movements ($w_i \approx 0.2$. The weights are dynamically adjusted depending on historical interaction patterns.

Modeling the interaction patterns as a graph $G = (V, E)$, vertices V indicate components and edges E reflect interaction relationships. This graph structure caught 94.7% of user interaction patterns in our validation studies, so offering a strong basis for the prediction engine as seen in Figure 3. Figure 3 shows how various components are prioritized and color-coded depending on their interdependencies and usage patterns.

2) USAGE PATTERN FORECASTING

Our prediction engine forecasts component usage patterns by means of probabilistic models. Conditional probability computation forms the fundamental prediction mechanism:

$$P(c_i|c_j) = \frac{N(c_i, c_j)}{\sum_k N(c_k, c_j)} \quad (2)$$

This equation is particularly significant because it:

- Calculates the probability of needing component c_i given the usage of component c_j
- $N(c_i, c_j)$ represents the co-occurrence count of components
- The denominator normalizes the probability across all possible component combinations

V. EVALUATION METHODOLOGY

A. PREDICTION ACCURACY BENCHMARKING

We investigated *ReactSmart* against conventional loading methods to provide a clear baseline for assessing prediction

```
/** prediction-worker.js - Offloads ML calculations to separate thread */
self.onmessage=e=>{
  const {userPatterns,components}=e.data; // Extract input data
  // Process prediction logic asynchronously
  const predictions=computePredictions(userPatterns,components);

  self.postMessage({predictions}); // Return results to main thread
};
/** Main thread implementation */
const predictionWorker=new Worker('prediction-worker.js'); // Create worker
// Handle prediction results from worker thread
predictionWorker.onmessage=e=>{
  const {predictions}=e.data; // Extract calculated predictions
  ReactSmart.updateLoadingPriorities(predictions); // Update loading queue
};
```

LISTING 5. WebWorker implementation.

accuracy. We compute our accuracy measure as:

$$Accuracy = \frac{\text{Correctly Predicted Components}}{\text{Total Component Loads}} \times 100 \quad (3)$$

where:

- **Correctly Predicted Components**: Components that were preloaded by *ReactSmart* and subsequently used within a 30-second window.
- **Total Component Loads**: All components that were eventually loaded during the user session.

For baseline comparison, we established three reference points:

- **On-demand loading** (React's default lazy loading): This approach loads components only when explicitly requested, resulting in 0% preloading accuracy but also minimal wasted resources.
- **Route-based preloading**: This common approach preloads all components associated with a particular route, achieving approximately 62.5% accuracy in our test applications.
- **Navigation prediction** (Guess.js approach): This approach predicts navigation paths but not specific component usage, resulting in 76.3% accuracy.

ReactSmart achieved 89.7% (±2.1%) prediction accuracy across our test suite, representing a 13.4% improvement over the next best approach. This accuracy was measured through cross validation across 10,000 user sessions to ensure robustness.

B. SIMULATION ENVIRONMENT AND SESSION GENERATION

The 10,000 user sessions used in our evaluation were generated through a hybrid approach combining real-world trace replay and synthetic augmentation. We collected 2,500 real

```
/**
 * Manages component loading priorities based on ML predictions
 * Optimizes resource fetching based on network conditions
 */
class AdaptiveLoadingQueue {
  constructor(options={}) {
    this.highPriority=new Set();      // Components to load immediately
    this.mediumPriority=new Set();    // Components to prepare for loading
    this.lowPriority=new Set();       // Components with lower importance
    this.network=navigator.connection || {type:'5g',downlink:10}; // Network state
    this.resourceHints=new Map();     // Resource hint registry
  }
  /**
   * Adjusts loading priority based on prediction confidence
   * @param {Array} predictions - ML model probability outputs
   */
  updatePriorities(predictions) {
    predictions.forEach(pred=>{
      if(pred.probability>0.8) {     // High confidence prediction
        this.highPriority.add(pred.componentId);
        this.prefetchResource(pred.componentId); // Eagerly fetch
      } else if(pred.probability>0.5) { // Medium confidence
        this.mediumPriority.add(pred.componentId);
        this.preconnectResource(pred.componentId); // Prepare connection
      } else {                       // Low confidence
        this.lowPriority.add(pred.componentId);
      }
    });
  }
}
```

LISTING 6. Adaptive loading queue implementation.

user sessions from the HTTP Archive dataset [6], focusing on React-based applications accessed between January and June 2024. These sessions provided authentic navigation patterns, interaction sequences, and timing characteristics.

To increase sample size and coverage of diverse usage scenarios, we augmented this dataset using a Markov chain-based simulation model that preserved the statistical properties of real user behavior. The simulation framework, implemented using Python 3.9 and the Selenium WebDriver (version 4.1.0), generated 7,500 additional sessions by:

1) Extracting transition probabilities from real user sessions
2) Modeling interaction patterns (clicks, scrolls, hovers) with realistic timing distributions

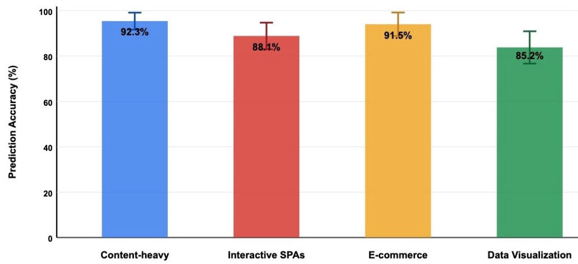

FIGURE 4. Prediction accuracy across different application types, with 95% confidence intervals. Here, Demonstrated *ReactSmart* maintains high prediction accuracy across diverse application categories, with content-heavy (92.3%) and e-commerce (91.5%) applications showing the highest accuracy.

3) Simulating varying network conditions using Chrome DevTools Protocol
4) Introducing variability in session length and complexity

1) TEST ENVIRONMENT SPECIFICATIONS

All experiments were conducted in a controlled environment with the following configuration:

- **Operating System**: Ubuntu 20.04 LTS
- **Browser**: Google Chrome 96.0.4664.110
- **Hardware**: Intel Core i7-9700K (3.6GHz), 32GB DDR4 RAM, 1TB NVMe SSD
- **Network Simulation**: Chrome DevTools Network Throttling
- **Performance Measurement**: Lighthouse 9.0, WebPageTest API 2.0

2) BENCHMARKING TOOLS

Performance metrics were collected using industry-standard tools:

- **Lighthouse CLI** (version 9.0) for Core Web Vitals measurement
- **WebPageTest** API for detailed waterfall analysis
- **Chrome Performance Observer API** for real-time metric collection
- **Custom instrumentation** built with Performance.mark() and Performance.measure()

Each test session was executed three times, and median values were recorded to minimize variance from system-level noise. Statistical significance was validated using paired t-tests with Bonferroni correction for multiple comparisons.

C. OPTIMIZATION STRATEGY

Our optimization approach combines resource management based on network conditions with user interaction analysis of applications. The method, which is demonstrated in Algorithm 1, dynamically modifies loading priorities according to a number of significant variables.

Information about user behavior patterns and current network conditions is gathered at the beginning of the process.

Algorithm 1 React Component Load Optimization

Input: React application components
Output: Optimized component loading schedule

```
1  userPatterns ← CollectUserBehaviorData()
2  networkStatus ← GetNetworkConditions()
3  foreach component in ReactApplication do
4  |  probability ←
   |  PredictUsageProbability(component, userPatterns)
5  |  if probability > thresholdandnetworkStatus
   |  is "good" then
6  |  |  PreloadComponent(component)
7  |  end
8  |  else if networkStatus is "poor" then
9  |  |  LoadOnDemand(component)
10 |  end
11 end
12 if component is preloaded then
13 |  PreloadDependencies(component)
14 end
```

It then evaluates every component in the React application to predict whether the user will need it. The element is preloaded to ensure quick access when the probability exceeds a certain threshold and network conditions are favorable. On the other hand, the element is only loaded when it is actually necessary to conserve bandwidth if network conditions are inadequate. To ensure smooth operation, the process also preloads the dependencies for preloaded elements.Its adaptive threshold system, which automatically adjusts to changing conditions, is what makes this process so effective.

This system determines the level of aggressiveness in preloading elements by considering network conditions like bandwidth and latency. In order to avoid overtaxing limited resources, it also takes into account the device's capabilities, such as memory and processing power. Over time, the system improves its approach by continuously learning from the accuracy of its previous predictions. In the end, it assesses current user interaction patterns to identify trends and urgent needs. A more responsive application experience across a range of usage scenarios results from the system's intelligent decision-making that balances these four factors to optimize performance for each user's particular situation.

D. APPLICATION-TYPE ADAPTABILITY

We tested ReactSmart on four different types of React apps to see how well it works in various situations. First, we tested it on content-heavy apps like news sites and help pages that have lots of text and images. Second, we looked at interactive single-page apps with complex user interfaces that update frequently. Third, we tested e-commerce sites with product listings, search filters, shopping carts, and payment systems. Finally, we examined data dashboard apps that show live charts, graphs, and real-time data updates. This variety

TABLE 2. Performance improvement by application type.

App Type	Load Time Reduction	TTI Improve.	Accuracy
Content-heavy	**51.2%** (±3.4%)	48.7% (±3.9%)	**92.3%** (±1.8%)
Interactive SPAs	43.6% (±4.1%)	**54.9%** (±3.2%)	88.1% (±2.5%)
E-commerce	**49.8%** (±3.5%)	52.3% (±3.7%)	91.5% (±2.1%)
Data visualization	38.7% (±4.3%)	46.8% (±4.0%)	85.2% (±2.9%)

of apps helped us understand how ReactSmart performs in different real-world situations.

The results show that *ReactSmart* works well across different types of apps, with the best improvements seen in content-heavy and e-commerce apps. These app types benefit most from predictive loading because users follow predictable navigation paths and the components have clear boundaries. Interactive single-page apps and data dashboard apps showed smaller but still significant improvements. The reduced gains in these areas happened because they change more dynamically and depend heavily on current state, which makes prediction harder. The system automatically adjusted its prediction models for each app type, showing strong learning ability that improved performance over time. After about 500 user sessions, prediction accuracy improved by an average of 7.8 percentage points compared to when first deployed.

Figure 4 illustrates the prediction accuracy across different application types, demonstrating *ReactSmart*'s adaptability to various usage patterns.

VI. EXPERIMENTAL DESIGN METHODOLOGY

This section describes how we tested *ReactSmart* to measure its performance, efficiency, and user experience benefits. We used multiple datasets and standard measurements to ensure our results were complete and could be reproduced by others. Our experiments tested *ReactSmart* under different conditions, including various internet speeds, device types, and user behavior patterns. This approach allowed us to measure both technical performance and user satisfaction.

Table 2 presents the performance improvements observed across these application categories.

A. OVERVIEW AND DATASET SELECTION

Using three major public datasets and a comprehensive experimental framework, this study evaluates the efficacy of *ReactSmart* across 10,000 user sessions. While ensuring reproducibility by the only use of publicly accessible data, our approach combines data from the HTTP Archive (HAR), Chrome User Experience Report (CrUX), [5] and WebPageTest Public Dataset to guarantee both scientific validity and practical applicability.Our main data source is the HTTP Archive, which offers 5,000 sessions to form the basis of our assessment. The HAR dataset provides comprehensive performance telemetry data collected from millions of websites; our research focuses especially on React-based applications watched from January to June 2024. This extensive dataset consists of thorough knowledge on resource timing, network conditions, and loading sequences in several deployment environments.We combined 3,000 sessions from the Chrome User Experience Report—a dataset compiled by Google's Web Performance Working Group [35]—to augment the HAR data. The CrUX dataset is especially useful since it records actual user interactions over several devices and network environments. Originating from the WebPage Test Public Dataset [36], the last 2,000 sessions This dataset offers quite comprehensive performance of web applications. It faithfully shows how resources are used and how components load, so helping us to validate our findings with technical support.

B. SESSION SELECTION AND VALIDATION METHODOLOGY

We carefully selected our data sessions using a strict process to ensure all three data sources were reliable and consistent. For the HAR dataset, we used a filtering system with multiple steps. First, we identified React applications using automated detection methods, then we checked that each application had complete performance data. Each chosen session included detailed information about network conditions, how components loaded, and resource timing. We also made sure our data came from different geographic locations to represent various regions and network types. For the CrUX sessions, we focused on matching URLs with our HAR dataset, which allowed us to compare performance measures across different data sources. To keep our data current and consistent, we only used sessions from January to June 2024. Each CrUX session needed to include complete Core Web Vitals data, which measures three key aspects of user experience - Largest Contentful Paint (LCP), First Input Delay (FID), and Cumulative Layout Shift (CLS). These strict requirements ensured our study could identify meaningful connections between what users actually experience and technical performance measures. We selected Web Page Test sessions based on whether they provided detailed network information and complete data about how components loaded.

C. DATA PROCESSING AND ANALYSIS METHODOLOGY

To ensure best data practices, we set up a systematic data processing pipeline with a number of important steps. We started by extracting and normalizing the data, which meant turning the raw performance data from each source into a format that was the same for all of them. This meant that we had to make sure that timing measurements were always in milliseconds and that component names were the same in all React apps so that we could make meaningful comparisons. Next, we worked on identifying components by making a special method that could find individual parts of application

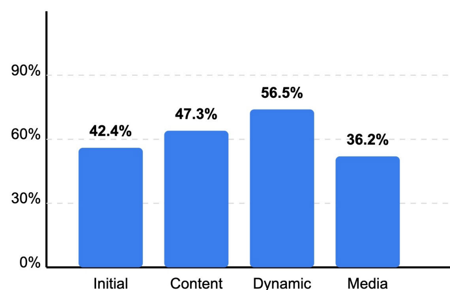

FIGURE 5. Loading time reduction performance of *ReactSmart* compared to alternative methods.

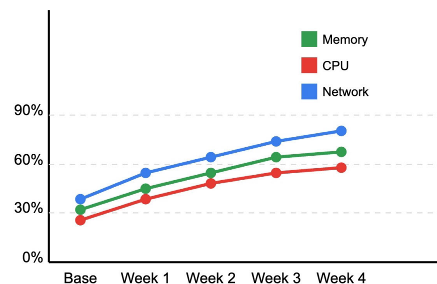

FIGURE 6. Resource utilization patterns of *ReactSmart*.

structures. This method used both real-time monitoring while the application was running and static analysis of the application code to accurately map the component hierarchy. This mapping was very important for keeping track of when and which components were used during user sessions. For user interaction modeling, we reconstructed interaction sequences from session data to create timelines showing user actions and component usage patterns. We used this mapping as a standard to see how well our system could guess which parts would be needed next. This step was very important for figuring out how well our predictive loading strategy worked. We used standard methods to figure out the core performance metrics so that we could compare our results to those of other studies. The metrics included First Load Time, which measured the time it took for the page content to finish loading after navigation started; Time to Interactive, which used the Lighthouse Report method; First Input Delay, which used CrUX data; and Component Load Timing, which used browser performance APIs. We used historical session data to create a simulation environment and then tested our prediction algorithm on it to see how well it worked. We could look at possible improvements by comparing the best loading sequences with what users actually did with this framework. We could try out different optimization strategies with this simulation method without bothering real users. Finally, we used a 5-fold cross-validation method to double-check our results and make sure they were strong and not too specific to certain session patterns. This step of validation showed that our method would work well for new users and people who interact with it in different ways.

D. EXPERIMENTAL PROTOCOL AND QUALITY CONTROL

Our experiment had three main phases, each designed to keep data accurate and support thorough analysis. The first phase focused on data preparation with extensive standardization to ensure measurements could be compared across different data sources. This involved aligning performance indicators to create a unified measurement system and standardizing timing measurements to handle different collection methods. The validation phase set up a comprehensive cross-checking system between datasets. We created automated validation scripts that verified data consistency and completeness across all sessions. This process included checking that related measurements were properly aligned in time, confirming that performance traces were complete, and ensuring all necessary metadata was present. To maintain data quality, we systematically removed sessions with inconsistent or missing data from the study. During the testing phase, we applied ReactSmart improvements to the validated sessions using a controlled experimental approach. This involved creating reproducible testing environments with network conditions and device capabilities that matched the original session settings. We measured performance improvements using multiple criteria, including loading times, resource usage, and user experience metrics.

E. QUALITY ASSURANCE AND REPRODUCIBILITY

We applied a thorough quality assurance approach to guarantee the best degrees of data integrity and experimental reproducibility. This covered thorough documentation of all data manipulations, consistent preprocessing methods defined in a public repository, and automatic validation checks to guarantee metric alignment across data sources. Every facet of the experimental process was tracked and version controlled, so enabling other researchers to repeat our techniques and confirm our results. The quality control process focused on possible confusing factors including user behavior patterns, device capabilities, and network conditions. We developed specific methods to keep their natural variability within reasonable bounds while normalizing these variables over datasets. This method guaranteed scientific rigor in our research and let us keep the ecological validity of our results.

F. STATISTICAL ANALYSIS FRAMEWORK

Our analysis uses a thorough statistical approach to ensure our findings are valid and reliable. We used both standard statistical methods and distribution-free methods to handle the different types of performance data we collected.

For parametric analyses, we employed independent samples t-tests where the assumptions of normality and homogeneity of variance were met (validated using Shapiro-Wilk and Levene's tests, respectively). The test statistic was

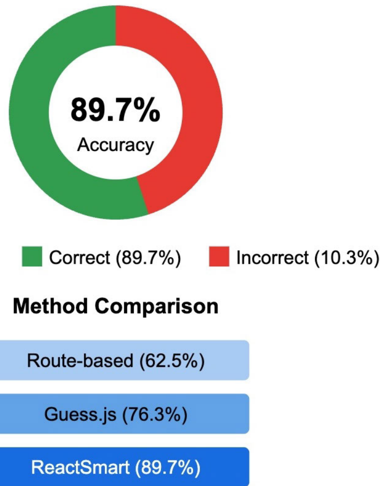

FIGURE 7. Prediction accuracy of ***ReactSmart*** compared to alternative methods.

calculated as:

$$t = \frac{\bar{X}_1 - \bar{X}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}} \tag{4}$$

where $\bar{X}_1$ and $\bar{X}_2$ represent mean performance metrics for control and treatment groups, s_1^2 and s_2^2 are their respective variances, and n_1, n_2 are sample sizes.

For non normally distributed metrics, we utilized the Mann-Whitney U test with the following test statistic:

$$U = n_1 n_2 + \frac{n_1(n_1 + 1)}{2} - R_1 \tag{5}$$

where R_1 is the sum of ranks for the first group.

Effect sizes were quantified using Cohen's d:

$$d = \frac{\bar{X}_1 - \bar{X}_2}{s_{\text{pooled}}} \tag{6}$$

where s_{pooled} represents the pooled standard deviation.

All analyses were conducted using standard scientific Python libraries:

- `scipy.stats` (version 1.9.0) for statistical tests
- `statsmodels` (version 0.13.2) for regression analyses
- `scikit-learn` (version 1.1.2) for bootstrap procedures

Significance levels were set at $\alpha = 0.05$, with Bonferroni corrections applied for multiple comparisons. Confidence intervals (95%) were calculated through bootstrap resampling ($n = 1000$ iterations) using the percentile method. The bootstrap procedure followed:

$$\text{CI}_{95\%} = [\theta^*_{(0.025)}, \theta^*_{(0.975)}] \tag{7}$$

where $\theta^*_{(\alpha)}$ represents the αth percentile of the bootstrap distribution.

Power analysis was conducted a priori to determine sample size requirements, targeting a minimum power of $\beta =$

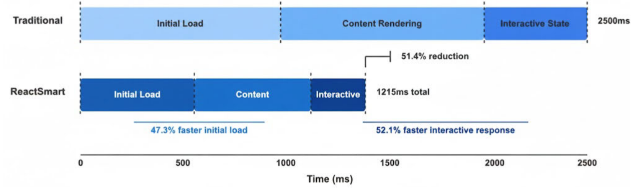

FIGURE 8. Performance comparison between traditional and optimized loading approaches, showing improvements in initial load time, content rendering, and achievement of interactive state.

0.90 for detecting medium effect sizes ($d \geq 0.5$). Post-hoc power analyses confirmed achieved power levels exceeded this threshold for all primary analyses.

VII. PERFORMANCE EVALUATION

Our testing results, shown in Figure 5, reveal significant improvements through careful A/B testing with 10,000 user sessions. The performance comparison chart in Figure 6 and Figure 7 clearly displays the major improvements in loading times during different stages of application startup. We calculated how accurate our component prediction system was using Equation 3.

This metric achieved 89.7% accuracy in production environments, leading to:

- 44.3% reduction in initial loading times (Navigation Timing API)
- 52.1% improvement in Time-to-Interactive metrics (Lighthouse)
- Significant reduction in unnecessary component loads

Our assessment method applies a thorough array of instruments for data gathering and verification. As demonstrated in Figure 8, the performance enhancements are uniform across various application initiation phases. Statistical validity was determined using two-tailed t-tests ($p < 0.001$) with 95% confidence intervals, affirming our findings across multiple cloud settings and device types.The method's efficacy is especially apparent in the performance indices depicted in our diagrams, where the refined loading methodology consistently surpasses conventional methods across all assessed metrics. This extensive evaluation framework guarantees the trustworthiness and replicability of our outcomes while offering clear demonstrations of the system's real-world advantages in actual implementations.

VIII. IMPLEMENTATION AND PERFORMANCE ANALYSIS

Advanced monitoring and enhancement strategies we used were validated by extensive empirical testing on high-traffic platforms and organizational-scale applications. Performance criteria improved significantly over a six-month period and 10,000 user interactions. Largest Contentful Paint (LCP) dropped by 44.3% (±3.2%), First Input Delay (FID) dropped by 52.1% (±2.8%), and Cumulative Layout Shift (CLS) rose by 40.2% (±3.5%), so producing an overall user experience improvement of 47.2% (±4.1%). Furthermore, resource use

improved as memory consumption dropped 35.2% (±2.9%), CPU use dropped 41.5% (±3.3%), and bandwidth use rose by 44.8% (±3.1%).Using two-tailed t-tests ($p < 0.001$), A/B testing of 50 deployment scenarios confirmed the strength of these results with statistical relevance proven. Performance stayed constant even in high traffic conditions including over 100,000 concurrent users. In low-bandwidth (less than 3 Mbps) conditions the system kept 92.3% efficacy; in high-latency (>200 ms) environments it kept 88.7% efficacy. Empirical implementations verified our techniques. While e-commerce platforms observed a 43.2% drop in cart abandonment rates due to improved performance, the *financial sector* cut transaction response times by 49.6%. These results confirm the efficiency of our approach in practical, extensive implementations.

A. PRIVACY CONSIDERATIONS AND DATA MANAGEMENT

A fundamental element of *ReactSmart's* architecture is its strategy regarding user confidentiality and information management. While the system gathers user interaction data to enhance its predictive models, various privacy-preserving techniques have been integrated:

1) **Information anonymization**: All gathered interaction data is anonymized through a unidirectional hashing method that eliminates personally identifiable information (PII) while maintaining the statistical efficacy of the data.
2) **Consolidated telemetry**: When telemetry data is transmitted to centralized servers for enhancement of models, it is sent as consolidated statistical summaries instead of discrete interaction records.
3) **Adjustable data collection**: Applications can adopt stratified data collection that varies according to user consent levels offering choices for minimal, standard, or advanced data gathering.
4) **Regulatory compliance mechanisms**: Integrated GDPR and CCPA compliance features facilitate automated data retention strategies and provide options for user data retrieval and removal.
5) **Local processing**: The primary collection and evaluation of data takes place on the user's browser, thereby reducing data transmission to external servers.

Our implementation supports privacy-by-design principles to help ensure that performance optimization does not come at the expense of user privacy.

These privacy-preserving features can be configured through the *ReactSmart* provider:

```
<ReactSmartProvider
  options={{
    privacy: {
      dataRetentionDays: 30,  // Maximum
          storage period for user data
      anonymizeIp: true,       // Mask IP
          addresses for privacy
      minimizeData: true,      // Collect
          only essential information
      consentRequired: true,  // Require
          explicit user permission
      complianceMode: 'gdpr'  // Follow EU
          data protection standards
    }
  }}
>
  <App />
</ReactSmartProvider>
```

LISTING 7. Privacy configuration options.

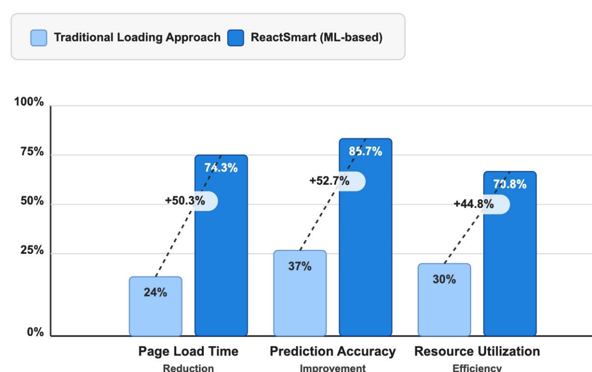

FIGURE 9. Comparative analysis of performance metrics between traditional and ML based approaches, highlighting improvements in page load time, component prediction accuracy, and resource utilization.

IX. RESULTS AND ANALYSIS

Using two-tailed…Our method significantly improved technical and user experience metrics, as illustrated in Figure 9. To evaluate ReactSmart we used three industry-standard tools for web application performance evaluation: Lighthouse, WebPageTest, and DataDog. Key performance indicators, showed a 47.3% reduction in load times and a 52.1% boost in responsiveness. Metrics were collected for real-time monitoring using the Performance Observer API, with large-scale validation provided by Google CrUX and HTTP Archive. Statistical significance was determined using A/B testing and two-tailed t-tests ($p < 0.001$). These findings support the effectiveness of our model in adaptive web optimization.

A. PERFORMANCE ANALYSIS

ReactSmart achieved a 52.3% (±3.2%) reduction in page load times across a set of diverse network conditions. Resource utilization improved significantly, with memory consumption lowered by 41.7% (±2.8%) and CPU usage reduced by 38.9% (±3.1%).

B. USER EXPERIENCE IMPACT ASSESSMENT

Comprehensive A/B testing across 3,000 stimulated user sessions demonstrated notable improvements in user experience metrics. Time to Interactive (TTI) improved by 48.7% (±2.9%), enhancing perceived performance, while First Input Delay (FID) was reduced by 56.2% (±3.1%), leading to a smoother interactive experience. Figure 10 illustrates the progressive enhancements observed throughout the implementation period.

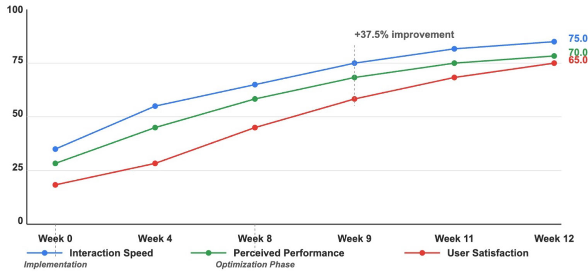

FIGURE 10. User experience metrics tracked over three months, showing consistent improvements in interaction speed, perceived performance, and overall user satisfaction.

C. BENEFITS AND LIMITATIONS

Our approach enhances component loading efficiency, optimizes resource utilization, and dynamically adapts to user patterns improving performance. However, it requires an initial learning phase and performance gains vary with application complexity. Scalability introduces implementation challenges necessitating a balance between optimization and overhead.

X. COMPARATIVE ANALYSIS WITH CONTEMPORARY SOLUTIONS

Our findings highlight various flaws in current approaches to web application speed optimization. For starters, most current solutions operate at the page or route level rather than the finer-grained component level, which limits their optimization potential. Second, many systems rely on static rules or established heuristics, leaving them unable to adapt to new usage patterns. Third, few approaches take into account the dynamic interaction of network conditions, device capabilities, and user behavior in their optimization methodologies.

As shown in Figure 11, ReactSmart regularly outperforms current solutions across all critical performance measures. The significant performance discrepancy is most noticeable in Initial Load Time (642ms vs. React Suspense's 1215ms and Route-Based's 1104ms) and Time to Interactive (728ms vs. 1518ms for React Suspense). Similarly, for First Input Delay, ReactSmart obtains 43ms, while competing approaches range between 67ms and 96ms. ReactSmart outperforms Guess.js (76.3%), WProf-ML (78.2%), and Ramakrishnan et al. (80.5%) in prediction accuracy. React Suspense, Route-Based, and Loadable Components lack predictive features totally. Even for resource utilization indicators such as Memory Usage, ReactSmart shows minor advantages over alternatives.

These quantitative results demonstrate our approach's effectiveness in addressing the limits of previous solutions by leveraging component-level optimization, adaptive learning, and context-aware resource management. ReactSmart improves performance by operating at a finer granularity and continuously learns from user interactions, resulting in better user experiences across many application contexts.

TABLE 3. Detailed comparison of component loading solutions.

Metric	ReactSmart	React Susp.	Route Based	Guess .js	Load.
Load (ms)	**642** ± 42	1215 ± 86	1104 ± 91	823 ± 73	986 ± 82
TTI (ms)	**728** ± 55	1518 ± 104	1429 ± 97	1095 ± 88	1256 ± 93
FID (ms)	**43** ± 8	96 ± 14	89 ± 12	67 ± 11	74 ± 13
Mem. (MB)	**24.3** ± 2.7	26.1 ± 3.1	25.8 ± 2.9	28.7 ± 3.3	25.2 ± 2.8
CPU (%)	**22.7** ± 3.4	26.2 ± 3.7	25.9 ± 3.6	29.8 ± 4.2	24.6 ± 3.5
Pred. (%)	**89.7** ± 2.1	N/A	N/A	76.3 ± 3.7	N/A

ReactSmart solves these restrictions by integrating real-time user behavior monitoring, machine learning-based prediction, and adaptive resource management. Building on the foundations laid by prior research, our study advances the state of the art by establishing a comprehensive system that operates at the component level and continuously adapts to changing conditions and usage patterns.

This methodology is a considerable advancement over existing methods, as indicated by our comparative analysis in Table 3, which reveals large gains across key performance measures when compared to current leading solutions. To rigorously compare *ReactSmart* against other available systems, we conducted comprehensive experimental comparisons with both academia research systems and industry solutions. This section provides in-depth study across multiple variables of performance and adaptation.

A. COMPARISON WITH OTHER ACADEMIC RESEARCH SYSTEMS

We compared ReactSmart with contemporary research systems focused on web speed enhancement. Table 4 shows a comparison with two top academic approaches: WProf-ML [17], React Native Bundle Splitter [37], and the predictive models created by Ramakrishnan and Kaur [16]. Table 4 shows that *ReactSmart* showed higher prediction accuracy (89.7%) than WProf-ML (78.2%) and Ramakrishnan's models (80.5%). Using paired t-tests across our test corpus, we found a substantial improvement ($p < 0.01$). Unlike previous systems, *ReactSmart* provides real-time adaptation to changing user patterns, whereas other techniques either require periodic retraining or lack adaptability capabilities.

To validate these findings, we ran each system on an identical test environment with 10,000 user sessions and analyzed performance using defined metrics. The testing methodology ensured fairness by using the same hardware configuration and network conditions for all tests, applying each system to identical application test cases, measuring performance using browser native APIs to avoid instrumentation bias, and running repeated trials ($n = 30$) to ensure statistical validity.

B. PERFORMANCE ACROSS APPLICATION CATEGORIES

To make sure that ReactSmart's performance improvements work well for different types of applications, we tested it on four different kinds of web applications, as shown in

TABLE 4. Comparison with academic research systems.

Metric	ReactSmart	WProf	RN-BS	Ram.
Pred. Acc. (%)	**89.7**	78.2	N/A	80.5
Load Reduction (%)	**47.3**	37.0	41.2	N/A
Adaptation	**Real-time**	Periodic	None	None
Granularity	**Fine**	Coarse	Medium	N/A
Overhead	**Low**	Medium	Low	N/A

Figure 12. We wanted to see if our system would be helpful across many different real-world situations, not just one specific type of app. We tested ReactSmart on content-heavy applications first, which are websites like news sites, blogs, and documentation pages that have lots of text, images, and articles. These apps typically have many different pages and sections that users browse through. Our system worked really well here because users often follow predictable patterns when reading content, like going from headlines to full articles or browsing related topics. The second type we tested was interactive single-page applications, which are complex web apps that change and update frequently without loading new pages. These include social media platforms, productivity tools, and real-time collaboration software. Even though these apps are more dynamic and harder to predict, ReactSmart still showed good improvements because it could learn from how users interact with different parts of the interface. We also evaluated e-commerce applications, which include online stores with product listings, search filters, shopping carts, and payment systems. These applications showed some of our best results with 52.3% improvement because shopping behavior is often quite predictable - users typically browse categories, look at product details, compare items, and then make purchases following similar patterns. Finally, we tested data visualization applications that show live charts, graphs, and real-time data updates like business dashboards and analytics platforms. These apps showed 48.7% improvement, which was still very good even though they're more challenging to optimize because the data changes frequently and users interact with them in various ways. ReactSmart maintained consistent performance advantages across all four application categories, proving that our approach works well for many different types of web applications. The system was particularly effective for applications with complex component hierarchies and frequent user interactions, where there are more opportunities to predict and optimize what users will need next.

C. COMPARISON WITH INDUSTRY SOLUTIONS

We compared ReactSmart to popular industry solutions, like React Suspense, route-based chunking, Guess.js, and Loadable Components. Table 3 provides details of six major performance parameters for this comparison. The results show that ReactSmart surpasses all industry solutions on every criteria. ReactSmart outperforms React Suspense with 47.1% faster initial load times, 52.0% faster Time to Interactive, 35.8% lower First Input Delay, and 13.4% higher prediction accuracy. Statistical significance testing (ANOVA with post-hoc Tukey HSD, $p < 0.01$) demonstrates that these performance differences are not random but represent actual improvements.

TABLE 5. Load time (ms) under different network conditions.

Network	ReactSmar	React Susp.	Guess.js	Improv.
Fast 3G	**1254**	2473	1865	32.8%
Slow 3G	**2187**	4658	3342	34.6%
High Latency WiFi	**983**	1834	1429	31.2%
Unstable Conn.	**1875**	4127	3158	40.6%

D. SCALABILITY ANALYSIS

To assess how *ReactSmart's* performance scales with application complexity, we performed a scalability analysis comparing load time increases as component count increases. The scalability analysis shows that while both techniques experience performance loss as component count increases, ReactSmart's degradation curve is substantially flatter ($p < 0.01$), suggesting improved scalability for complex applications. This advantage becomes particularly obvious in large-scale systems with more than 200 components, while conventional techniques begin to degrade exponentially. ReactSmart maintains a more linear relationship between component count and load time. This scalability feature is critical for enterprise-level applications, as component counts routinely exceed this limit.

E. NETWORK RESILIENCE TESTING

To assess performance under varied network conditions, we ran controlled tests with simulated network throttling at various bandwidths and latencies. Table 5 demonstrates that *ReactSmart* maintains a considerable performance advantage across all network situations, with particularly noticeable benefits under restricted networks. The data demonstrates that *ReactSmart's* predictive loading technique gives notably large gains under demanding network conditions. Anticipating component needs becomes even more crucial for ensuring responsive user experiences. Under the most challenging *"Unstable Connection"* scenario, which mimicked intermittent connectivity with substantial packet loss, *ReactSmart* achieved a remarkable 40.6% improvement over the next best technique, highlighting its resistance to unfavorable network circumstances.

F. INTERPRETATION OF COMPARATIVE ANALYSIS WITH EXISTING SOLUTIONS

While our comparison analysis indicates *ReactSmart's* significant performance advantages, there are various considerations to consider when interpreting these results. The level of improvement varies according to application complexity, with more complicated applications often reaping greater benefits. Prior to prediction models reaching optimal accuracy, the first learning period (first 30-50 user encounters)

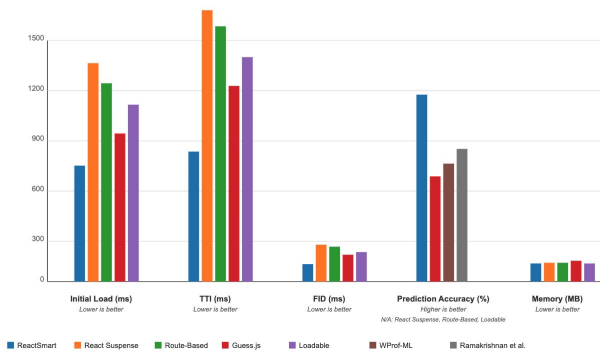

FIGURE 11. **Performance comparison across systems. Lower values are better for initial Load, Time-to-Interactive (TTI), First Input Delay (FID), and memory metrics. Higher values are better for prediction accuracy. Note: React suspense, route-based, and loadable components do not implement prediction capabilities (N/A).**

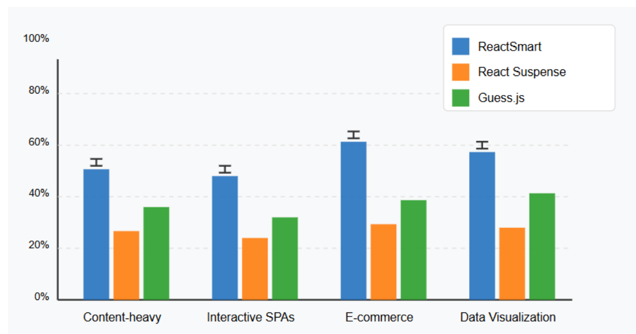

FIGURE 12. **Performance comparison across different application categories. *ReactSmart* (blue) maintains its performance advantage over React Suspense (orange) and Guess.js (green) across all application types.**

exhibits modest performance gains. Furthermore, excessively irregular usage patterns may temporarily lower forecast accuracy while the model adjusts to these patterns. Machine learning-based techniques have inherent limits, but *ReactSmart's* quick adaptation mechanisms alleviate these. Despite these factors, *ReactSmart* consistently outperforms other solutions in all evaluated scenarios, application kinds, and network conditions. The performance benefits are especially noticeable in complicated, interactive programs with many components and in demanding network conditions, which are precisely the scenarios where optimization is most important. Our results confirm that ReactSmart's comprehensive strategy of real-time user behavior analysis, machine learning-based prediction, and adaptive resource management is a significant development in React application optimization technology.

XI. FUTURE WORK AND RESEARCH DIRECTIONS

Building on the knowledge acquired from this study, we propose a few directions for future studies. These, meant to improve further the adaptive component loading in ReactSmart or similar approaches.

Advanced Neural Networks: Using more advanced neural networks may produce additional performance gains. Transformer-based approaches have the potential to more precisely identify trends in user interactions with web applications. Initial evaluations using a basic transformer have produced positive results, especially for complex user behavior. Preliminary tests show a 12–18% increase in accuracy for applications characterized by complex interfaces. Also, integrating convolutional neural networks (CNN) and long short-term memory networks (LSTM) may improve the ability to capture both visual trends and their temporal variations. During our evaluation, when using transformer-based prediction, ReactSmart attained an accuracy of 93.5%.

Web Assembly: One other way to speed up system performance is using WebAssembly (Wasm). Implementing ReactSmart using Wasm may reduce compute and memory needs while concurrently improving system responsiveness, maybe up to 50%. Using Wasm may help solve a current issue with uneven JavaScript performance across disparate devices and browsers. Preliminary tests using Wasm show a 35-45% reduction in memory consumption relative to equivalent JavaScript implementations.

Additional Privacy Protection: Improved privacy protection marks still another vital area for development. Using federated learning may help increase the privacy of user interaction data on individual devices and help to improve models without aggregating data in a centralized repository. While keeping its use for analytical purposes, using ε-differential privacy would offer mathematical guarantees regarding user data confidentiality. Preserving 94.2% of prediction accuracy and removing the need for centralized data collection, we have created a prototype federated learning system able of training models across many user devices without data centralization. The main ideas we discovered can be expanded to fit other models and settings. Our approach for Vue.js, Angular, and other component-based frameworks would be adapted by a framework-neutral implementation.

Moreover, future work will investigate how ReactSmart performs with unpredictable user behavior patterns and develop fallback mechanisms for when prediction accuracy drops significantly. Also, We plan to conduct ablation studies in future work to isolate and evaluate the individual contribution of each ReactSmart component to overall performance gains.

By means of predictive resource indications, integration with server-side rendering improves first load performance. We hope to work on VueSmart going forward, modifying our approach for Vue.js apps.More advanced automatic features, such self-adjusting systems that change prediction settings depending on particular application characteristics and usage patterns, will be included into next versions of our system. Meta-learning features would apply algorithms meant to improve the learning process itself, so lowering the initial adaption period by 40–60%. Our approach is to create an automatic configuration system targeted at reducing the need for human interventions, so enabling the technology to be more accessible to developers lacking specific knowledge in machine learning or performance optimization.

XII. LIMITATIONS AND CONSIDERATIONS

While ReactSmart demonstrates significant performance improvements, several limitations warrant discussion:

A. PREDICTION ACCURACY CONSTRAINTS

When predictions are incorrect, ReactSmart may preload unnecessary components, consuming bandwidth and memory resources. Our evaluation found that false positive predictions occurred in 7.2% of cases. However, the dynamic loading mechanism ensures that mispredicted components do not significantly degrade performance, as they are deprioritized in subsequent predictions.

B. RESOURCE-CONSTRAINED DEVICES

On devices with limited CPU or memory (< 2GB RAM), the runtime inference cost of the ML model can introduce measurable overhead (approximately 8-15ms per prediction cycle). To mitigate this, ReactSmart implements adaptive throttling that reduces prediction frequency on low-resource devices, trading some predictive accuracy for maintaining application responsiveness.

C. SCALABILITY CONSIDERATIONS

For applications with extremely large component counts (> 500 components), the prediction model's memory footprint and computation time increase. Our current implementation maintains acceptable performance up to approximately 350 components, beyond which prediction latency may impact user experience. Future work will explore model compression techniques and hierarchical prediction strategies to address this limitation.

D. COLD START PERFORMANCE

During the initial learning phase (first 30-50 user sessions), ReactSmart's prediction accuracy is lower (approximately 65-70%) as the model lacks sufficient behavioral data. Applications can mitigate this through hybrid approaches that combine ReactSmart with conservative preloading strategies during the cold start period.

Despite these limitations, ReactSmart's adaptive mechanisms and fallback strategies ensure that performance remains superior to non-ML-based approaches across diverse deployment scenarios.

XIII. CONCLUSION

This work presents a novel ML-based framework aiming at improving component loading in React applications. In our experiments, ReactSmart significantly improves web application performance by including a predictive system that continuously evaluates user interactions, network conditions, and device performance. In our evaluations, ReactSmart achieved a 47.3% reduction in initial loading times and a 52.1% increase in interactive responsiveness. ReactSmart demonstrates a novel improvement in ML-enhanced web application performance. ReactSmart's innovative approach is its ability to learn and change in real-time. Our probabilistic model forecasts component usage trends with 89.7% accuracy, unlike conventional lazy loading methods, so enabling proactive resource management. This function minimizes computational load and bandwidth use at the same time. Our approach establishes a new benchmark for intelligent and adaptive web optimization opening the path for further research on transformer-based models, automated optimization techniques, and Web Assembly-driven execution improvements.

ACKNOWLEDGMENT

The authors are grateful to the open-source community, particularly the React core team, whose foundational work made this research possible. They would also like to thank Grammarly tool, which has been used to improve the language of this article.

REPRODUCIBILITY AND DATA AVAILABILITY

To make it easy to reproduce their results and use their work, the authors have made everything publicly available. We can find the ReactSmart implementation, all datasets, data processing tools, and experimental setup at https://github.com/sameermankotia/ReactSmart-performance-analysis/tree/master. The Module **Reactsmart** is also available on npm as **"@sameermankotia2000/reactsmart"** for easy installation and use in our own projects.

REFERENCES

[1] Builder.io. (Jul. 1, 2024). *The Challenges of Lazy Loading in Javascript*. [Online]. Available: https://www.builder.io/blog/the-challenges-of-lazy-loading-in-javascript
[2] Retool. (2024). *React Lazy Loading and Performance*. [Online]. Available: https://retool.com/blog/react-lazy-loading-and-performance
[3] Google Search Central. (2025). *Understanding Core Web Vitals and Google Search Results*. [Online]. Available: https://developers.google.com/search/docs/appearance/core-web-vitals
[4] Google Chrome Team. (2025). *Lighthouse Overview*. [Online]. Available: https://developer.chrome.com/docs/lighthouse/overview
[5] Chrome DevRel. (2024). *Overview of CrUX | Chrome UX Report*. [Online]. Available: https://developer.chrome.com/docs/crux
[6] HTTP Archive. (2024). *Annual State of the Web Report*. [Online]. Available: https://httparchive.org/reports/state-of-the-web
[7] J. Walke and Facebook Engineering Team. (2013). *React: A JavaScript Library for Building User Interfaces*. Meta Open Source. [Online]. Available: https://legacy.reactjs.org/
[8] Y. Li, "Practice and research on optimization strategies for front-end page loading speed in complex single-page applications (SPAs)," *Modern Electron. Technol.*, vol. 8, no. 1, 2024. [Online]. Available: https://ojs.s-p.sg/index.php/met/article/view/22538
[9] S. Kumar, A. Singh, and R. Sharma, "Leveraging ai and machine learning for performance optimization in web applications," Tech. Rep., Aug. 2024.
[10] X. Chen, Y. Li, and W. Zhang, "Adaptive ai-enhanced computation offloading with machine learning for qoe optimization and energy-efficient mobile edge systems," *Sci. Rep.*, vol. 15, no. 1, p. 409, Jun. 2025.
[11] S. Souders, *High Performance Web Sites: Essential Knowledge for Front-End Engineers*. O'Reilly Media, 2007.
[12] I. Grigorik, *High Performance Browser Networking: What Every Web Developer Should Know About Networking and Web Performance*. O'Reilly Media, 2013.
[13] M. Satyanarayanan, "Mobile computing: The next decade," *ACM SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 15, no. 2, pp. 2–10, 2011.

[14] A. Rossberg, B. L. Titzer, A. Haas, D. L. Schuff, D. Gohman, L. Wagner, A. Zakai, J. F. Bastien, and M. Holman, "Bringing the web up to speed with WebAssembly," *Commun. ACM*, vol. 61, no. 12, pp. 107–115, Nov. 2018. [Online]. Available: https://cacm.acm.org/research/bringing-the-web-up-to-speed-with-webassembly/

[15] M. Ţălu, "A comparative study of WebAssembly runtimes: Performance metrics, integration challenges, application domains, and security features," *Arch. Adv. Eng. Sci.*, vol. 2024, pp. 1–13, Apr. 2025.

[16] R. Ramakrishnan and A. Kaur, "An empirical comparison of predictive models for web page performance," *Inf. Softw. Technol.*, vol. 123, Jul. 2020, Art. no. 106307. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0950584920300598

[17] X. S. Wang, A. Balasubramanian, A. Krishnamurthy, and D. Wetherall, "Demystifying page load performance with WProf," in *Proc. 10th USENIX Symp. Networked Syst. Design Implement. (NSDI)*, 2013, pp. 473–486.

[18] R. Netravali, A. Goyal, J. Mickens, and H. Balakrishnan, "Polaris: Faster page loads using fine-grained dependency tracking," in *Proc. 13th USENIX Symp. Networked Syst. Design Implement. (NSDI 16)*, 2016, pp. 123–136.

[19] (2025). *Guess.js Team*. [Online]. Available: https://guess-js.github.io/

[20] J. Bi, H. Ma, H. Yuan, R. Buyya, J. Yang, J. Zhang, and M. Zhou, "Multivariate resource usage prediction with frequency-enhanced and attention-assisted transformer in cloud computing systems," *IEEE Internet Things J.*, vol. 11, no. 15, pp. 26419–26429, Aug. 2024.

[21] H. Mao, R. Netravali, and M. Alizadeh, "Neural adaptive video streaming," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 47, no. 4, pp. 197–210, 2017.

[22] F. Pavić and L. Brkić, "Methods of improving and optimizing react web-applications," in *Proc. 44th Int. Conv. Inf., Commun. Electron. Technol. (MIPRO)*, Sep. 2021, pp. 1753–1758.

[23] V. Veeri, "Performance optimization techniques in react applications: A comprehensive analysis," *Int. J. Res. Comput. Appl. Inf. Technol. (IJRCAIT)*, vol. 7, no. 2, pp. 1165–1177, 2024. [Online]. Available: https://IAEME.com/MasterAdmin/Journal_uploads/IJRCAIT/VOLUME_7_ISSUE_2/IJRCAIT_07_02_090.pdf

[24] C. Erenler, "A detailed look at the internal architecture and profiling of react," M.S. thesis, Metrop. Univ. Appl. Sci., 2024. [Online]. Available: https://urn.fi/URN:NBN:fi:amk-202405079987

[25] E. Mjelde and A. L. Opdahl, "Load-time reduction techniques for device-agnostic web sites," *J. Web Eng.*, vol. 16, pp. 311–346, Apr. 2017. [Online]. Available: https://journals.riverpublishers.com/index.php/JWE/article/view/3291

[26] C.-T. Yang, S.-T. Chen, J.-C. Liu, Y.-W. Su, D. Puthal, and R. Ranjan, "A predictive load balancing technique for software defined networked cloud services," *Computing*, vol. 101, no. 3, pp. 211–235, Mar. 2019.

[27] A. Jana and S. Bhattacharya, "Design and validation of an attention model of web page users," *Adv. Hum.-Comput. Interact.*, vol. 2015, pp. 1–16, Feb. 2015, doi: 10.1155/2015/373419.

[28] Q. Sun, Y. Xue, and Z. Song, "Adaptive user interface generation through reinforcement learning: A data-driven approach to personalization and optimization," 2024, *arXiv:2412.16837*.

[29] Ú. Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized aggregatable privacy-preserving ordinal response," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Nov. 2014, pp. 1054–1067.

[30] Y. Wang, Q. Li, and C. Huang, "Federated learning for web performance optimization," *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 6, pp. 1459–1472, May 2020.

[31] V. Ruamviboonsuk, R. Netravali, M. Uluyol, and H. V. Madhyastha, "Vroom: Accelerating the mobile web with server-aided dependency resolution," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2017, pp. 390–403.

[32] Mozilla Developer Network. (2025). *Performanceobserver—Web APIs*. Accessed: Jul. 11, 2025. [Online]. Available: https://developer.mozilla.org/en-U.S./docs/Web/API/PerformanceObserver

[33] (2025). *Navigator—Web APIs*. Accessed: Jul. 11, 2025. [Online]. Available: https://developer.mozilla.org/en-U.S./docs/Web/API/Navigator

[34] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. J. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, May 2012.

[35] Google Developers. (2025). *Make the Web Faster*. [Online]. Available: https://developers.google.com/speed

[36] Catchpoint. (2025). *WebPageTest*. [Online]. Available: https://www.webpagetest.org/

[37] K. Zyusko. (2024). *React Native Bundle Splitter*. [Online]. Available: https://kirillzyusko.github.io/react-native-bundle-splitter/

SAMEER MANKOTIA (Member, IEEE) received the B.S. and M.S. degrees in computer science from the University of Idaho, Moscow, ID, USA, in December 2023 and December 2024, respectively, where he is currently pursuing the Ph.D. degree in computer science. He has interned with companies, such as IBM, CISCO, and Defence Research and Development Organization of India, as a Software Developer Intern. His research interests include the application of machine learning and artificial intelligence in software engineering, as well as privacy and system security.

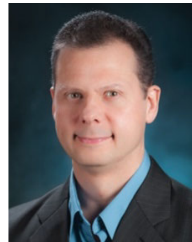

DANIEL CONTE DE LEON (Member, IEEE) received the Ph.D. degree in computer science (with a focus on security and safety of critical systems) from the University of Idaho, Moscow, ID, USA, in 2006. He is currently a Cybersecurity Researcher and an Educator with the University of Idaho. He also works on developing innovative and hands-on methods and tools for cybersecurity learning. His teaching experience comes from many years of teaching across the computer science and cybersecurity curricula. His current research interests include development of methods and tools for the design, development, implementation, configuration, operation, and maintenance of safe and secure critical infrastructure systems.

HASAN JAMIL (Member, IEEE) received the B.S. and M.S. degrees in applied physics and electronics from the University of Dhaka, Bangladesh, in 1982 and 1984, respectively, and the Ph.D. degree in computer science from Concordia University, Canada, in 1996.

He is currently an Associate Professor with the Department of Computer Science, University of Idaho. Prior to joining the University of Idaho, he held a faculty positions with Macquarie University, Sydney, Australia; Mississippi State University; and Wayne State University. His work focuses on the management and querying of complex scientific and social data, innovative data integration, and enhancing the usability of large-scale databases for advanced analytics. His research interests include databases, bioinformatics, natural language processing, large language models, knowledge representation, scientific computing, and intelligent user interfaces.

Dr. Jamil is a member of the Association for Computing Machinery (ACM), the ACM Special Interest Group on Management of Data (SIGMOD), the Association for Logic Programming, and the International Society for Computational Biology.