



# What makes Ethereum blockchain transactions be processed fast or slow? An empirical study

Michael Pacheco<sup>1</sup> · Gustavo A. Oliva<sup>1</sup> · Gopi Krishnan Rajbahadur<sup>2</sup> · Ahmed E. Hassan<sup>1</sup>

Accepted: 19 December 2022 / Published online: 3 February 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

## Abstract

The Ethereum platform allows developers to implement and deploy applications called DApps onto the blockchain for public use through the use of smart contracts. To execute code within a smart contract, a paid transaction must be issued towards one of the functions that are exposed in the interface of a contract. However, such a transaction is only processed once one of the miners in the peer-to-peer network selects it, adds it to a block, and appends that block to the blockchain. This creates a delay between transaction submission and code execution. It is crucial for DApp developers to be able to precisely estimate when transactions will be processed, since this allows them to define and provide a certain Quality of Service (QoS) level (e.g., 95% of the transactions processed within 1 minute). However, the impact that different factors have on these times have not yet been studied. Processing time estimation services are used by DApp developers to achieve predefined QoS. Yet, these services offer minimal insights into what factors impact processing times. Considering the vast amount of data that surrounds the Ethereum blockchain, changes in processing times are hard for DApp developers to predict, making it difficult to maintain said QoS. In our study, we build random forest models to understand the factors that are associated with transaction processing times. We engineer several features that capture blockchain internal factors, as well as gas pricing behaviors of transaction issuers. By interpreting our models, we conclude

---

Communicated by: Philipp Leitner

✉ Michael Pacheco  
mpacheco@cs.queensu.ca

Gustavo A. Oliva  
gustavo@cs.queensu.ca

Gopi Krishnan Rajbahadur  
gopi.krishnan.rajbahadur1@huawei.com

Ahmed E. Hassan  
ahmed@cs.queensu.ca

<sup>1</sup> Software Analysis and Intelligence Lab (SAIL), School of Computing Queen's University, Kingston, ON, Canada

<sup>2</sup> Centre for Software Excellence (CSE), Huawei, Kingston, ON, Canada

that features surrounding gas pricing behaviors are very strongly associated with transaction processing times. Based on our empirical results, we provide DApp developers with concrete insights that can help them provide and maintain high levels of QoS.

**Keywords** Transaction processing time · Smart contracts · Ethereum · Blockchain · Machine learning · Regression model · Model interpretation

## 1 Introduction

A blockchain provides a secure and decentralized infrastructure for the execution and record-keeping of digital transactions. A *programmable* blockchain platform supports *smart contracts*, which are stateful, general-purpose computer programs that enable the development of *blockchain-powered applications*. Ethereum is currently the most popular programmable blockchain platform. In Ethereum, these blockchain-powered applications are known as decentralized applications (DApps).

Transactions are the means through which one interacts with Ethereum. Ethereum defines two types of transactions: *regular* transactions and *contract* transactions. Regular transactions enable transfers of cryptocurrency (Ether, or simply ETH) between end-user accounts. In turn, contract transactions invoke functions defined in the interface of a smart contract. When an end-user triggers a certain functionality on the frontend of a DApp (e.g., checkout the items in the shopping cart of an online shopping application), such a request is translated into one or more blockchain contract transactions. DApps themselves facilitate these transactions by submitting them to the blockchain (e.g., Ethereum). That is, end-users are completely unaware that a blockchain is in use, as transactions are handled entirely by the DApp.

In Ethereum, all transactions must be paid for, including contract transactions. However, transaction prices are not prespecified. Transaction issuers need to decide how much they wish to pay for each transaction by the assignment of the *gas price* parameter, paid in Ether (ETH). The higher one pays for a transaction, the faster it is generally processed. This is because there is a financial incentive that *mining nodes* (i.e., those who choose, prioritize, and process all transactions on Ethereum) receive for transactions that they process and verify. This incentive is a function of the gas prices that are associated with the transactions they choose to process. As a result, the execution of code on programmable blockchains such as Ethereum takes time, as the corresponding transaction(s) must first be chosen by the mining node that will append the next block to the blockchain. Consistently setting up transactions with high gas prices is not an option, since it would render the DApp economically unviable. Hence, for a DApp to be profitable, it needs to issue transactions with the lowest gas price that fulfills a predefined Quality of Service (QoS) (e.g., transactions processed within 1 minute in 95% of the cases). To achieve such a goal, DApp developers make use of *processing time estimation services*.

Although existing processing time estimation services are useful, they offer little to no insight into what drives transaction processing times in Ethereum (except for gas prices). This happens either because the underlying machine learning model that is used by these services is a complete black-box (e.g., Etherscan's Gas Tracker<sup>1</sup>) or because the service

---

<sup>1</sup><https://etherscan.io/gastracker>

does not clearly indicate the features that drive the model's decisions (e.g., EthGasStation<sup>2</sup>). Ultimately, the processing time estimations made by these services impact the gas price choices made by DApp developers for their transactions. We thus argue that these services should be *interpretable*. Without explanations or insights into the underlying models, the estimations provided by these services lack reliability, robustness, and trustworthiness. In practice, this means DApp developers struggle to define proper gas prices and manage transaction submission workloads (e.g., decide the best time to submit a given set of transactions). An in-depth investigation of the performance of these popular estimation services can be found in our prior work (Pacheco et al. 2022).

Building an explainable transaction processing time model entails discovering and understanding the factors that influence (or are at least correlated with) processing times. As a first step towards addressing this challenge, in this study we set out to determine and analyze the factors that are associated with transaction processing times in Ethereum. In the following, we list our research questions and the key results that we obtained:

**RQ1: How well do blockchain internal factors characterize transaction processing times?** To answer this RQ, we first build a random forest model with an extensive set of features which capture internal factors of the blockchain. (e.g., the difficulty to mine a block, or the characteristics of the contracts and transactions) of the Ethereum blockchain. These features are derived directly from the contextual, behavioral, and historical characteristics associated with 1.8M transactions. Next, we interpret the model using adjusted  $R^2$  and SHAP values.

*Our model achieves an adjusted  $R^2$  of 0.16, indicating that blockchain internal factors explain little of the variance in processing times.*

**RQ2: How well does gas pricing behavior characterize transaction processing times?** Similar to the stock market, the Ethereum blockchain ecosystem is impacted by a plethora of external factors, including speculation, the global geopolitical situation, and cryptocurrency regulations. Many of these factors are largely unpredictable and difficult to be captured in the form of engineered features. Yet, all these factors are known to impact the *gas pricing behavior* of a large portion of transaction issuers. Hence, in RQ2 we build a model that aims to account for the gas pricing behavior of transaction issuers (i.e., the mass behavior). We do so by engineering additional features that indicate how the price of a given transaction compare to that of recently submitted transactions.

*Our model achieves an adjusted  $R^2$  score of 0.53 (2.3x higher than the baseline model), indicating that it explains a substantially larger amount of the variance in processing times. The most important feature of our model is the median percentage of transactions processed per block, in the previous 120 blocks (~30min ago), with a gas price below that of the current transaction.*

**Paper organization** The remainder of this paper is organized as follows. Section 2 defines the key concepts employed throughout this study. Section 3 provides a motivating example that emphasizes the importance having explainable transaction processing time predictions. Section 4 describes our study methodology, including the data collection, our engineered feature-set, and the model construction. Sections 5 and 6 show the motivation, approach, and findings associated with the research questions addressed in this paper. Section 8 discusses related work. Section 9 discusses the threats to the validity of our study. Finally, Section 10 states our concluding remarks.

---

<sup>2</sup><https://ethgasstation.info>

## 2 Background

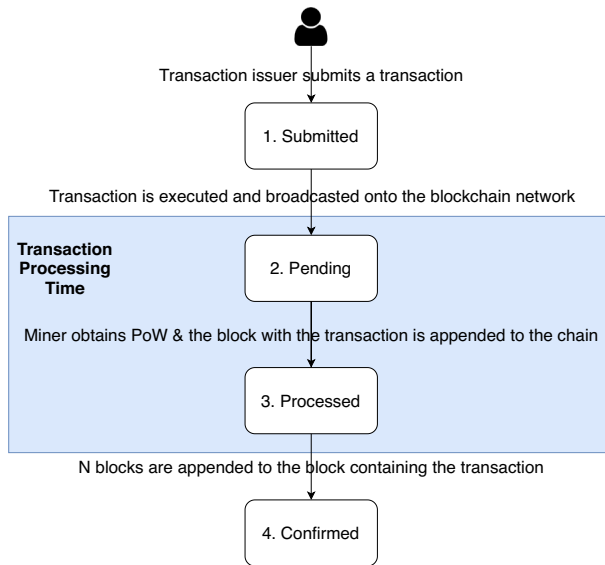
In this section, we define the key concepts of the Ethereum blockchain that are employed throughout our study.

**Blockchain** A blockchain is commonly conceptualized as a ledger which records all transaction-related activity that occurs between peers on the network. This activity data is stored through the usage of blocks, which is a data structure that holds a set of transactions. Each of these transactions can be identified by an identifier called the transaction hash. Blocks are appended to one another to form a linked list type structure. A block stored in the blockchain becomes immutable once a specific number of blocks has been appended to that block. This is because for one to change the information within a specific block, they must also change all blocks appended after it as well.

**Transactions** Transactions allow peers to interact with each other in the context of a blockchain. There are two types of transactions used in the Ethereum blockchain: *user transactions* and *contract transactions*. User transactions are employed to transfer cryptocurrency (known as Ether in Ethereum) to another user, who is identified by a unique address. Contract transactions allow users (e.g., DApp developers) to execute functions defined in smart contracts. Smart contracts are immutable, general-purpose computer programs deployed onto the Ethereum blockchain.

**Gas and transaction fees** Transactions sent within the Ethereum blockchain require a transaction fee to be paid before being sent. This fee is calculated by  $gas\ usage \times gas\ price$ . The *gas usage* term refers to the amount of computing power consumed to process a transaction. This value is constant for user transactions, costing 21 GWEI, equivalent to  $2.1E-8$  ETH. For contract transactions, this value is calculated depending on the bytecode executed within a smart contract, as each instruction has a specific amount of gas units required for it to execute (Wood 2019). The *gas price* term defines the amount of Ether the sender is prepared to pay per unit of gas consumed to process their transaction. This parameter is set by the sender, and allows users to incentivize the processing of their transaction at a higher priority (Signer 2018), as miners which process the transaction are rewarded based on the transaction fee value.

**Transaction processing lifecycle** The lifecycle of a transaction transitions through various states, which are illustrated in Fig. 1. The first state of the transaction processing lifecycle is when a user first submits their transaction  $t$  to the Ethereum blockchain (*submitted* state). Next, this transaction  $t$  is broadcasted through the network by a peer-to-peer based communication method. Once the first node discovers transaction  $t$ , the transaction enters the *pending* state as it enters the pending pool of transactions. As time passes this transaction  $t$  is continuously discovered by an increasing amount of nodes deployed on the network, which track and record all transaction related activity, including those that are pending. Mining nodes will compete to win the Proof-of-Work (PoW) (Jakobsson and Juels 1999), which will allow the winner to append their new block  $b$  to the blockchain. If transaction  $t$  has been processed and recorded in block  $b$ , transaction  $t$  then transitions into the *processed* state (i.e., at the *block timestamp* of  $b$ ). Transaction  $t$  will transition into the next state, *confirmed*, once a specific number of blocks  $n$  has been appended after  $b$ . There is no universal agreement on what the value of  $n$  should be. For instance,  $n = 7$  is stated in the Ethereum whitepaper (Buterin 2014). Considering that blocks are appended every 15s on average,



**Fig. 1** Transaction lifecycle in the Ethereum blockchain

$n = 7$  translates to approximately 1m 45s. Alternatively, some cryptocurrency exchanges define  $n$  as a higher value. For instance, CoinCentral uses  $n = 250$ , which translates to approximately 1h 2m 30s (Comben 2018).

**Proof-of-work** Proof-of-Work (PoW) is a consensus protocol that requires nodes to solve a difficult mathematical puzzle. This protocol ensures that the strategy used to solve the puzzle is optimal (i.e. no solution for the puzzle exists other than a brute force solution). Fortunately the verification of such a solution is much cheaper in terms of computation. Ultimately this allows the nodes, or entities within the network who do not know the authority of one another, to trust each other and build a valid transaction ledger without requiring a centralized third-party authority (e.g., a bank).

**Next-generation DApps** Commonly, DApps expose the internal blockchain complexities to end-users, ultimately requiring that they install a *wallet* (e.g., Metamask). A wallet is a computer program (commonly a web browser extension) that enables one to submit transactions in a blockchain without having to set up their own node. However, submitting transactions with a wallet still requires end-users to set up *transaction parameters*. Such a task is error-prone by construction (Oliva and Hassan 2021), as not all end-users are familiar with blockchain concepts and their intricacies. Other problems include end-users getting confused with the user interface of wallets or even making typos while inputting the transaction parameter values. As an illustrative example, an Ethereum user called Proudbitcoiner spent 9,500 USD in transaction fees to send just 120 USD to a smart contract<sup>3</sup>. The user said the following (*gas limit* and *gas price* are two transaction parameters and 1 GWEI = 1E-9 ETH):

<sup>3</sup>The transaction logs can be seen at <https://etherscan.io/tx/0x5641c38bf30fd19ad3332100762017573ee676f35250dcf4a976bb0cfe31ac2f>

*“Metamask didn’t populate the ‘gas limit’ field with the correct amount in my previous transaction and that transaction failed, so I decided to change it manually in the next transaction (this one), but instead of typing 200000 in “gas limit” input field, I wrote it on the “Gas Price” input field, so I paid 200000 Gwei for this transaction and destroyed my life.”*

Due to incidents such as the one above, alongside end-user onboarding barriers, more recently DApp developers have started to engineer their applications using a new paradigm, which hides all blockchain complexities from end-users. In this new paradigm, the burden of setting up transaction parameters is on the DApp developers. DApp developers, in turn, can charge end-users with a flat fee (e.g., payable via credit card), such that they do not need to hold Ether or use wallets, ultimately easing their onboarding process. We refer to these DApps as the next-generation DApps (Oliva and Hassan 2021). Next-generation DApps are a reality and include DApps such as MyEtherWallet<sup>4</sup>.

### 3 Motivating Example

Assume that some DApp development organization (e.g., a company) wants to convert their web application into an Ethereum DApp. In order to ensure a high quality end-user experience, the company will submit contract transactions itself (instead of pushing this burden to end-users) (Oliva and Hassan 2021). Since end-users will be unhappy if transactions take too long to be processed, the company defines the following quality of service (QoS): 90% of all transactions shall be processed within 5 minutes. Considering that transactions cost ETH to be sent, this QoS would be easily achievable if the company’s budget were infinite. Unfortunately, this is not the case for any real-world DApp development organization. In a real-world scenario, we assume that organizations aim to provide such a QoS while *also* ensuring that the gas prices for their transactions are as low as possible. We refer to this balance as the *time-cost balance*.

Organizations can leverage processing time estimation services to help them achieve a proper time-cost balance. For example, assume a developer (John) is chosen to implement the code that is responsible for submitting the DApp’s transactions. Just before submitting a transaction, John requests a lookup table from his favorite estimation service (e.g., Eth-GasStation). A lookup table contains processing time predictions for every possible value of gas price. This table is calculated by the estimation service based on recently processed transactions. John then attempts to implement an algorithm that processes the lookup table and selects the  $\langle \text{gas price}, \text{estimated processing time} \rangle$  pair that most closely aligns with the predefined QoS (given the actual processing times of his previously submitted transactions). John quickly realizes that implementing such an algorithm is very hard and error-prone because *estimation services offer little to no explanation for their predictions*. For instance, assume that John needs to submit transactions shortly, but gas prices are high overall. Should John wait a little in the hope that prices will come down or should he submit the transactions right away? It is difficult to answer this question without knowing the explanations behind the predicted processing times. Assume that John knew that *network congestion level* was the key factor driving the predicted processing times shown in his current lookup table. In that case, John would be able to analyze the duration of past network congestions to make a more informed decision regarding the best time to submit his transactions. However, in

---

<sup>4</sup><https://www.myetherwallet.com/>

reality, John does not even know if one or more factors are influencing transaction processing times. As a result, it becomes nearly impossible for John to design more advanced algorithms for transaction pricing. John actually struggles to manage transaction submission workloads and meet the predefined QoS. Ultimately, the lack of explainable estimation services limits John's ability from finding the optimal time-cost balance for the company's DApp.

The aforementioned motivating example shows that there is a need for some unified, free, and open-source alternative to the existing estimation services that would provide DApp developers such as John with explainable transaction processing time predictions. As a first step towards addressing this need, in this paper we study the factors that are associated with transaction processing times in Ethereum.

## 4 Methodology

We build a machine learning model with a comprehensive feature-set in order to determine the factors that are strongly associated with the processing times of transactions in Ethereum. In this section, we explain our approaches to data collection (Section 4.1), feature engineering (Section 4.2), and model construction (Section 4.3).

### 4.1 Data Collection

In this section, we describe the data sources that we used (Section 4.1.1) and the approach that we followed to collect the data (Section 4.1.2).

#### 4.1.1 Data Sources

Our study relies on two data sources, namely Etherscan and Google BigQuery:

**Etherscan** Etherscan<sup>5</sup> is a popular real-time dashboard for the Ethereum blockchain platform. Etherscan has its own nodes in the blockchain, which are used to actively monitor the activity of the network in real time. This includes monitoring various details regarding transactions, blocks, and smart contracts. We use the data from this website to obtain transaction processing times (our dependent feature) and to engineer several features used in our model (our independent variables, such as pending pool size and network utilization level). Data from Etherscan has been used in several blockchain empirical studies (Oliva et al. 2020; Oliveira et al. 2021; Pacheco et al. 2022; Pierro and Rocha 2019)).

**Google BigQuery** Google BigQuery<sup>6</sup> is an online data warehouse that supports querying through SQL. In particular, Google also actively maintains several public datasets, including an Ethereum dataset<sup>7</sup>. This dataset is updated daily and contains metadata for several blockchain elements, including transactions, blocks, and smart contracts. Most of the features that we engineer rely on data collected from this dataset.

---

<sup>5</sup><https://etherscan.io/>

<sup>6</sup><https://cloud.google.com/bigquery>

<sup>7</sup>The dataset is `bigquery-public-data.crypto.ethereum`

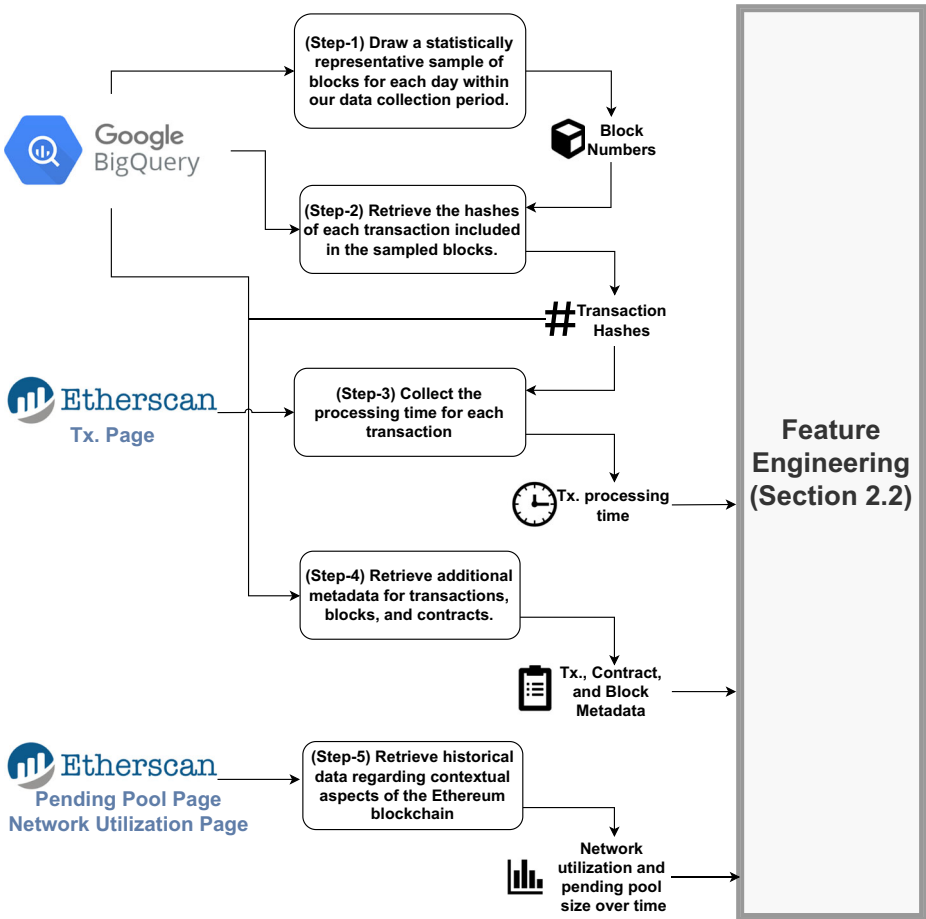


Fig. 2 An overview of our data collection approach. The dashed lines represent a connection to a data source

### 4.1.2 Approach

Figure 2 provides an overview of our data collection approach. In the following, we discuss each of the data collection steps.

**(Step-1) Draw a statistically representative sample of blocks for each day within our data collection period** To bring the data size to more manageable levels, we draw a statistically representative sample of mined blocks for each day within a one-month period. To draw the samples we use a setting of a 95% confidence level and  $\pm 5$  confidence interval, similarly to several prior studies (Tagra et al. 2022; Zhang et al. 2019; Boslaugh and Watters 2008). Following, for each day in our time period we calculate the statistically representative sample of blocks, and sample the corresponding amount. This approach results in a sample of 362 blocks per day on average, and a total of 10,865 blocks.

**(Step-2) Retrieve the hashes of each transaction included in the sampled blocks** We proceed by collecting and saving the *transaction hashes* of all transactions that have been included within the sampled blocks. The *transaction hash* is a unique identifier used to



differentiate between transactions. We use these transaction identifiers in several of the data collection steps to collect various associated metadata with each of these transactions. We collect the *transaction hash* for all 1,825,042 (1.8M) transactions included in the 10,865 sampled blocks based on the previous step.

**(Step-3) Collect the *processing time* for each transaction** Next we begin to collect the *processing time* provided by the Etherscan platform. For a given transaction  $t$ , the *processing time* is the estimated difference between the time that  $t$  has been included in a block that has been appended to the blockchain, and the time the user executed the submission of  $t$  onto the blockchain. We retrieve this data by using the *transaction hash* of each transaction in our set to access the corresponding *processing time* details provided by Etherscan (defined as *confirmation time* by Etherscan).

**(Step-4) Retrieve additional metadata for transactions, blocks, and contracts** We then leverage Google BigQuery to retrieve additional information regarding transactions, including information about the blocks and smart contracts they are associated with (when applicable). Examples of the types of the collected metadata from this source includes details such as gas price, block number, and network congestion. These metadata are obtained by querying the *transactions*, *blocks*, and *contracts* tables.

**(Step-5) Retrieve historical data regarding contextual aspects of the Ethereum blockchain** In addition to the *processing time*, the Etherscan platform provides historical data that describes some of the contextual information of the Ethereum blockchain. In particular, we collect data which describes: 1) the amount of transactions in the pending pool<sup>8</sup>, and 2) the network utilization<sup>9</sup>. We then map this set of data to each transaction, associating the relevant context of the blockchain to the time each transaction was executed.

## 4.2 Feature Engineering

To understand the features that impact the transaction processing times, we collect and engineer a variety of features. We explain the rationale behind focusing on these dimensions, as well as the rationale for designing each of the features in our set. Features were carefully crafted based on our scientific (Oliva et al. 2020; Kondo et al. 2020; Oliva and Hassan 2021; Zarir et al. 2021) and practical knowledge of Ethereum (e.g., buying/selling/trading ETH and watching the cryptocurrency market fluctuations). We list each of the features that we use for building our models in Tables 1 and 2.

**1. Features that capture internal factors** We first engineer features that aim to capture internal factors of the Ethereum blockchain (RQ1), which are summarized in Table 1. These features span across a few dimensions, including contextual, behavioral, and historical dimensions. We explain these dimensions below.

- *Contextual.* Contextual data describes the circumstances in which a given transaction is being processed. This information can be associated with the entire blockchain network, and also the transactions themselves. For example, the network utilization at the time a transaction was executed is a contextual feature. The network utilization should help provide indicators of changes to transaction processing times, such as increase in network traffic causing processing times to be delayed. Similarly, a feature which

<sup>8</sup><https://etherscan.io/chart/pendingtx>

<sup>9</sup><https://etherscan.io/chart/networkutilization>

**Table 1** The list of independent features used to capture internal factors of the Ethereum blockchain, along with their rationale. These features are used to build our models in Section 5 in order to predict transaction (tx) processing times. The ● symbol represents data retrieved from Etherscan, while the ★ symbol represents data retrieved from Google BigQuery. Braces ({} ) group related features

Type	Name	Description	Rationale
Contextual	contract_block_number★	The block number of the smart contract that the tx has interacted with.	Contracts that have been deployed onto the Ethereum blockchain long ago might impact processing times by containing inefficient code, security vulnerabilities, or bugs, causing them to be avoided by miners.
	contract_bytecode_length★	The bytecode length of the smart contract the tx has interacted with.	Larger smart contract sizes might result in an increase in execution time, which might also impact processing times.
	is_{erc721, erc20}★	A flag indicating whether the smart contract the tx has interacted with is an ERC721 or ERC20 (token-related) contract.	Contract txs. might be prioritized by miners, and might also have longer execution times, compared to user txs.
	to_contract★	A flag indicating whether the tx is interacting with a contract.	
	pending_pool●	The total amount of txs. in the pending pool that have yet to be processed, at the time closest to the time of sending the tx.	The amount of txs. in the pending pool, along with network utilization, reflects the network traffic at a given point in time. We predict an increase in these metrics should also result in increased tx processing times.
	net_util●	The network utilization of the network, as calculated by Etherscan.	
	day★	The day of the week the tx was executed on.	Contextual patterns of the network, such as increased usage, might revolve around certain days of the week or hours in the day.
	hour★	The hour of the day the tx was executed on.	

**Table 1** (continued)

Type	Name	Description	Rationale
Behavioral	gas_price_gwei*	The gas price value in GWEI set by the user of the tx.	The gas price is directly related to tx processing times - miners receive a monetary reward based on the gas prices of the txs. inside the blocks that they append to the chain. As a result, txs. with higher gas prices should have fast processing times.
	tx_nonce*	The total number of txs. sent by the user up until the current tx.	User inexperience might be associated with setting inadequate tx parameters, most importantly the gas price. Also, txs. with different nonce values might also be prioritized differently by various miners.
	value*	The value of ETH the user is sending along with their tx.	Differences in the ETH value might be prioritized differently by various miners.
	tx_gas_limit*	The gas limit value set by the user of the tx.	Difference in gas limits might be prioritized differently by various miners. For example, miners might avoid specific txs. if they think it might result in out of gas error.
	input_length*	The length of the input data associated with the tx.	Longer inputs might take longer to process in smart contracts, leading to longer execution time, and also longer processing time.
Historical	{total, avg, med, std}_txs_120*	The total, average, median, and standard deviation of all processed txs. across all previous 120 blocks, and the previous block.	Increases in the number of txs. being processed in the recent past is likely to increase processing times in the near future.
	total_txs_1*	The total number of txs. processed in the previous block.	

**Table 1** (continued)

Type	Name	Description	Rationale
	avg_gas_price_gwei_prev_day●	The average gas price of txs. processed the previous day.	
	{avg, med, std}_difficulty_120★	The average, median, and standard deviation of the difficulty of the previous 120 blocks.	The difficulty is directly related to the time taken to process blocks. This value is altered based on the time taken by miners to append blocks to the Ethereum blockchain.
	difficulty_1★	The difficulty of the previous block.	
	{avg, med, std}_func_gas_usage_120★	The average, median, and standard deviation of gas usage of the processed txs. in the previous 120 blocks, targeting the same smart contract function as the current tx.	The gas usage from the smart contract txs. are constantly targeting may indicate the popularity and utilization of such contract, which may be prioritized by miners in a particular way.
	{avg, med, std}_pending_pool_120●	The average, median, and standard deviation of the number of txs. in the pending pool per minute, during the appending of the previous 120 blocks.	As the amount of txs. in the pending pool increases, miners have more txs. to prioritize, which should increase processing times.
	{avg, med, std}_pend_prices ★	The average, median, standard deviation of gas prices (GWEI) of txs from the same user that are currently in the pending pool and not processed at the same time of the recorded tx For the first tx, the value of these features is zero.	The more txs a user has that are simultaneously waiting to be processed, the longer it could take for a single tx of that user to be processed due to the gas prices and nonce of the txs.
	num_pending★	The number of txs from the same user that are currently in the pending pool and not processed at the same time of the recorded tx.	

**Table 2** The list of independent features used to capture gas pricing behavior, along with their rationale. These features are used to build our models in Section 6 in order to predict transaction (tx) processing times. The ● symbol represents data retrieved from Etherscan, while the ★ symbol represents data retrieved from Google BigQuery. Braces ({} ) group a set of related features

Name	Description	Rationale
{avg, med, std}_num_{above, same, below}_120★	The average, median, and standard deviation of txs. processed per block, in the previous 120 blocks, with a gas price above, equal to, and below the current tx.	The various aggregations of txs. sent in previous blocks, combined with the aggregations of the associated gas price value, should hold stronger explanatory power than these gas price and processing time features alone. As a result, these features should be powerful in explaining variations in processing times.
{avg, med, std}_pct_{above, same, below}_120★	The average, median, and standard deviation of the percentage of txs. processed per block, in the previous 120 blocks, with a gas price above, equal to, and below the current tx.	
num_{above, same, below}_{1, 120}★	The total number of txs. processed per block, in the previous 120 blocks and previous block, with a gas price above, equal to, and below the current tx.	
pct_{above, same, below}_{1, 120}★	The percentage of the total number of txs. processed per block, in the previous 120 blocks and previous block, with a gas price above, equal to, and below the current tx.	
gas_price_cat_enc★	The label encoded gas price category of the tx calculated using the percentiles [0.2, 0.4, 0.6, 0.8, 1.0] of gas prices from all txs. from the previous 120 blocks.	The gas price is an incentive for miners to prioritize one tx over another. This feature aims to determine whether a certain gas price is <i>very low</i> , <i>low</i> , <i>medium</i> , <i>high</i> , or <i>very high</i> by comparing it to the prices of recent txs.
{avg, med, std}_gas_price_{1, 120}★	The average, median, and standard deviation of tx gas prices per block, in the previous 120 blocks and previous block	Such aggregations capture current gas prices being paid, and thus should help in predicting future trends in gas prices and its relationship to processing times.

**Table 2** (continued)

Name	Description	Rationale
past_{avg, med, std}.time●	The average, median, and standard deviation of the processing times for the most recent 100 txs. in the previous 120 blocks, with the most similar gas prices as the current tx.	The aggregations of processing times recorded for the most similar txs. within the previous blocks will give the model a sense of how long the current tx should be taking to be processed.
closest_tx_pr.time●	The processing time of the most recent tx processed in the previous 120 blocks with the most similar gas price as the current tx.	

defines the smart contract a transaction had interacted with in its lifetime also provides contextual insights.

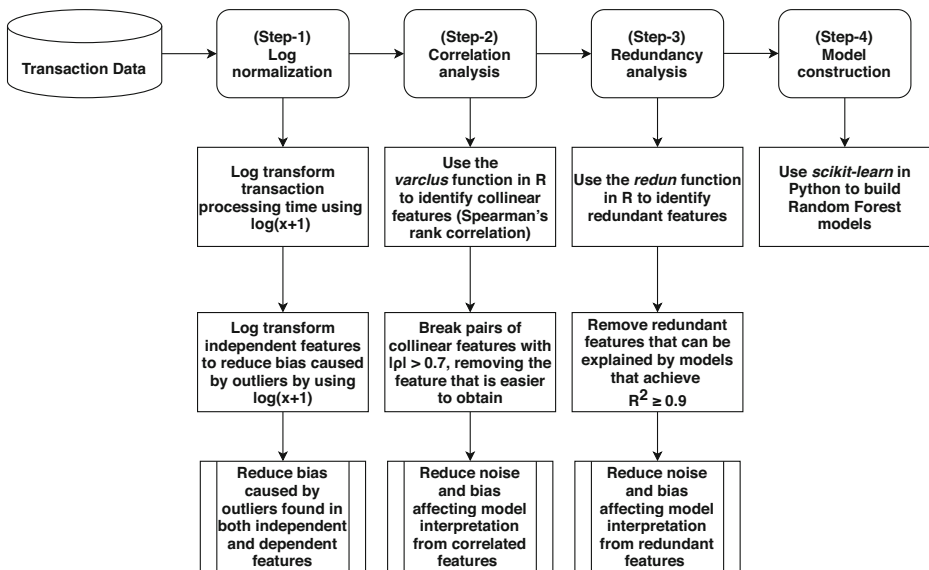
- *Behavioral*. Behavioral features include information which relate to the user and the particular values that were chosen for each of the transaction parameters. Depending on the experience of the user who is sending the transaction, the values set for these parameters might impact transaction processing times. Most notably, miners have the freedom to prioritize transactions to process using their own algorithms. However, it is common for miners to prioritize processing transactions with higher gas price values, as they receive a reward for processing a particular transaction based on this value. An inexperienced user who is not familiar with gas price values might set too low of a value, resulting in an extremely long processing time.
  - *Historical*. Historical features relate to features based on historical data from transactions that have already been processed in the blockchain. Such metrics are widely used in machine learning for predictive modelling tasks, as data of the past is used to predict outcomes of the future. For example, a feature which describes the average processing time for processed transactions in the previous  $x$  blocks can be used by models to understand and predict processing times for similar transactions in the future.
- 2. Features that capture changes in gas pricing behavior.** We also engineer features that aim to track pricing behavior over time (RQ2). Our features primarily involve comparing the gas prices of transactions processed in recent blocks with that of the current transaction at hand. These features and their associated rationale are summarized in Table 2.

### 4.3 Model Construction

In this section, we describe our model construction steps. The goal of our constructed model is to explain how features related to internal factors and features related to gas pricing behavior influence the transaction processing times respectively. We also aim to answer which of the studied features influence the transaction processing times the most. To do this, we design and follow a generalizable and extensible model construction pipeline to generate interpretable models, as outlined in prior studies (Lee et al. 2020; Mcintosh et al. 2016; Pacheco et al. 2022; Thongtanunam and Hassan 2020). Rather than focusing on generating models that are meant to be used to predict transaction processing times with high to perfect levels of accuracy, we instead leverage machine learning as a tool to understand the impact of these features. Our model construction pipeline is summarized in Fig. 3. We explain each step in more detail below.

**(Step-1) Log normalization.** We observe that the distribution of transaction processing time is right skewed. Therefore, we log transform the data with the function  $\log(x + 1)$  to reduce the skew. Furthermore, we also log transform all the independent features of the type *numeric* with an identical function to reduce the bias caused by outliers, similar to prior studies (Lee et al. 2020; Mcintosh et al. 2016; Menzies et al. 2007; Tantithamthavorn et al. 2016).

**(Step-2) Correlation analysis.** Several prior studies show that presence of correlated features in the used data to construct a machine learning model leads to unreliable and spurious insights (Midi et al. 2013; Mcintosh et al. 2016; Tantithamthavorn and Hassan 2018; Lee et al. 2020). Therefore, in this step we outline our process to remove correlated and redundant features in our data. To remove correlated features, we use the `varclus`



**Fig. 3** A summary of our model construction approach

function from the `Hmisc` package in R<sup>10</sup>, which generates a hierarchical clustering of the provided features, based on the correlation between them. We choose Spearman's rank correlation ( $\rho$ ) to determine the level of correlation between pairs, as rank correlation can assess both monotonic and non-monotonic relationships. Similar to prior studies (Lee et al. 2020; McIntosh et al. 2016; Tantithamthavorn and Hassan 2018), we use a threshold of  $|\rho| > 0.7$  to deem a pair of features as correlated. For each pair of correlated features, we choose the feature that is easier to collect or calculate in practice. We summarize our correlation analysis in Table 3 (Appendix A). As a result of such an analysis, we removed a total of 41 correlated features.

**(Step-3) Redundancy analysis** In addition to correlated features, redundant features may also negatively impact the generated explanations of a machine learning model (Midi et al. 2013; McIntosh et al. 2016; Tantithamthavorn and Hassan 2018; Lee et al. 2020). Hence, we choose to remove redundant features. To do so, we use the `redun` function from the `Hmisc` package in R<sup>11</sup>. This function builds several linear regression models, each of which selects one independent feature to use as the dependent feature. It then calculates how well the selected dependent features can be explained by the independent ones, through the  $R^2$ . It then drops the most well-explained features in a stepwise, descending order. This feature removal process is repeated until either: 1) none of the features can be predicted by the models resulting in an  $R^2$  greater than a set threshold, or 2) removing the feature causes another previously removed feature to be explained at a level less than the threshold. In our study, we use the default  $R^2$  threshold of 0.9. As a result of such an analysis, we removed a total of 1 redundant feature, namely `std_pct_below_120`.

<sup>10</sup><https://www.rdocumentation.org/packages/Hmisc/versions/4.5-0/topics/varclus>

<sup>11</sup><https://www.rdocumentation.org/packages/Hmisc/versions/4.5-0/topics/redun>



**(Step-4) Model construction** At this step, we fit our random forest model using the `RandomForestRegressor` function from the `scikit-learn` package in Python<sup>12</sup>. We choose a random forest model as it provides an optimal balance between predictive power and explainability. Such models are equipped to fit data better than simple regression models, potentially leading to more robust and reliable insights (Lyu et al. 2021). Although they are not inherently explainable, several feature attribution methods can and have been used to derive explanations using these models (Breiman 2001; Lundberg et al. 2018; Tantithamthavorn and Jiarapakdee 2021). Random forests have also been widely used in empirical software engineering research to derive insights about the data (Rajbahadur et al. 2021; Tantithamthavorn and Jiarapakdee 2021; Tantithamthavorn et al. 2020). Historically these models have shown good performance results for Software Engineering tasks (Yatish et al. 2019; Bao et al. 2021; Fan et al. 2020; Aniche et al. 2020). Finally, the specific random forest implementation that we choose is natively compatible with the `shap`<sup>13</sup> python package. The latter implements the SHAP interpretation technique, which we employ to address our research questions (Sections 5 and 6).

#### 4.4 Model Validation

We use the adjusted  $R^2$  to evaluate how well our models can predict transaction processing times. To compute the adjusted  $R^2$  of our constructed random forest model, we set up a model validation pipeline that splits data into train, test, and validation sets (Fig. 4). Our train dataset needs to be as large as possible, otherwise our model cannot identify possible seasonalities in the data (e.g., evaluate the influence of the *day of the week* in transaction processing times) nor changes in the blockchain context (e.g., network conditions are unlikely to change in small periods of time). We thus define the training set as the first 28 days of data, the validation set as the 29<sup>th</sup> day, and the test set as the final day. We split the data along its temporal nature to avoid data leakage (Lyu et al. 2021). We then use the training set to generate 100 bootstrap samples (i.e., samples with replacement and of the same length as the training set). For each bootstrap sample, we first fit a random forest model on that sample. Next, we use the validation set to hyperparameter tune that model as recommended by prior studies to ensure that our constructed model fits the data optimally. To do this we utilize a Random Search algorithm, and choose the `RandomizedSearchCV` (Bergstra and Bengio 2012) from `scikit-learn`<sup>14</sup>. We choose to hyperparameter tune the parameters: `max_depth`, `max_features`, and `n_estimators`. Once we discover the optimal hyperparameters, we retrain the model using the chosen parameters (on the same bootstrap sample) and evaluate it on the test set by adjusted  $R^2$  measure. While prior studies typically measure the explanatory power of a constructed model by measuring its  $R^2$  score on the training set, we instead measure the performance of our hyperparameter-tuned model on a hold-out test set. Our reasoning is that random forest models are by default generally constructed in a way which causes all training data points to be found in at least one terminal node of the forest. As a consequence, it is common for the random forest to predict correctly for the majority of instances of the train set, resulting in an  $R^2$  that is very close to 1 (typically in the 0.95 to 1.0 range) (Liaw 2010). In addition, we use *adjusted* version of the  $R^2$  measure, which mitigates the bias introduced by the high number of features used in

<sup>12</sup><https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestRegressor.html>

<sup>13</sup><https://shap.readthedocs.io/en/latest/index.html>

<sup>14</sup>[https://scikit-learn.org/stable/modules/generated/sklearn.model\\_selection.RandomizedSearchCV.html](https://scikit-learn.org/stable/modules/generated/sklearn.model_selection.RandomizedSearchCV.html)

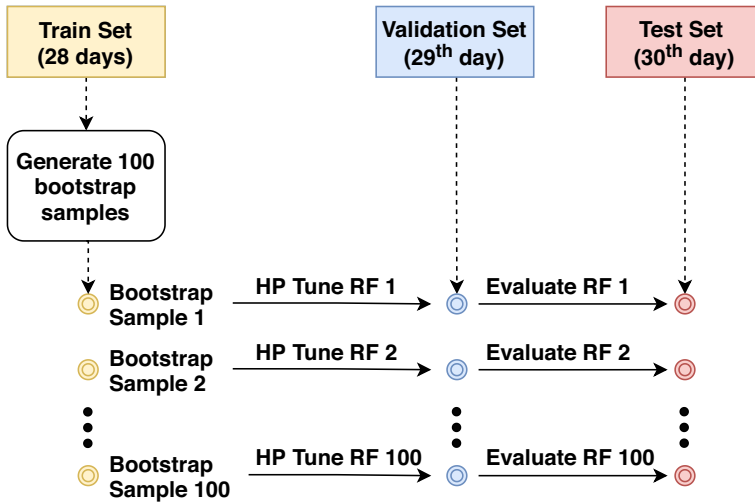


Fig. 4 Our model validation approach. RF stands for Random Forest and HP stands for hyperparameter

our model (Harrell 2015). As we generate 100 bootstrap samples, we obtain 100 adjusted  $R^2$  values. We report summary statistics for this distribution of adjusted  $R^2$  values in the research questions below.

**Summary**

- Total collected transactions: 1.8M transactions, spanning over 10.8k blocks.
- Data sources: Etherscan and Google BigQuery.
- Pieces of collected data: Metadata for blocks (e.g., number of included transactions), transactions (e.g., gas price parameters, processing times), contracts (e.g., whether it implements the ERC20 interface), network utilization over time, pending pool size over time.
- Raw feature-set (pre-filtering): 75 features (9 contextual, 5 behavioral, 20 historical, and 41 for gas pricing behavior)
- Final feature-set (post-filtering): 30 features (6 contextual, 4 behavioral, 9 historical, and 11 for gas pricing behavior).
- Machine learning algorithm: Random forest regressor.

## 5 RQ1: How Well do Blockchain Internal Factors Characterize Transaction Processing Times?

**Motivation** It is currently unknown which factors from the internal workings of the Ethereum blockchain best characterize transaction processing times. There exists an exhaustive set of information recorded on the blockchain at any given moment in time. For DApp developers, it is unclear which of these factors are useful in predicting processing times, and thus which should be monitored and analyzed.


**Approach** We analyze the model's explanatory power and feature importances. The details are shown below:

- *Analysis of a model's explanatory power.* To determine the explanatory power of blockchain factors, we engineer features based on the *internal* data recorded and available in the Ethereum blockchain (Table 1). Next, we use these feature to construct (Section 4.3) and validate (Section 4.4) our random forest model. As a result of applying our model validation approach, we obtain a distribution of 100 adjusted  $R^2$  scores. This score distribution determines how much of the total variability in processing times is explained by our random forest model, thus serving as a good proxy for the explanatory power of our model. We judge the explanatory power of our model by analyzing the mean and median adjusted  $R^2$  scores.
- *Analysis of feature importance.* To determine what individual features are most strongly associated with processing times, we compute and analyze feature importance using SHAP (Lundberg and Lee 2017; Molnar 2020). SHAP is a model-agnostic interpretation technique based on the game theoretically optimal *Shapley values*. SHAP first determines an average prediction value known as the *base rate*. Next, for each dataset instance, SHAP determines the Shapley value of each feature. The Shapley value informs the degree to which each feature moves the base rate (positively or negatively), in the same scale as the dependent feature. The final prediction is explained as:  $base\ rate + \sum_{i=1}^n Shapley(f_i)$ , where  $f_i$  is one of the  $n$  features of the model. The importance of a feature  $f_i$  corresponds to its distribution of Shapley values across all dataset instances. More details about SHAP can be found in Appendix B.

To compute Shapley values, we first train a single random forest model on the entire *train set* (28 consecutive days) using the procedures that are outlined in Section 4.3. Next, we use the `TreeExplainer` API from the `shap` python package to compute the Shapley value of each feature for each dataset instance. As a result, each feature will have an associated Shapley value distribution of size  $N$ , where  $N$  is the total number of rows in our train dataset. Next, we convert these Shapley value distributions into *absolute* Shapley value distributions. The rationale of such a conversion is that a feature  $f_1$  with a Shapley value of  $x$  is as important as a feature  $f_2$  with a Shapley value of  $-x$ . Finally, we employ Scott Knott algorithm to rank the distributions of absolute Shapley values. The feature ranked *first* is the one with the highest Shapley values. We note that there can be a rank tie between two or more features, which indicates that their underlying shapley distributions are indistinguishable from an effect-size standpoint. We also note that we use Menzies' implementation (Menzies 2020) of the original SK algorithm (Scott and Knott 1974), which (i) replaces the original parametric statistical tests with non-parametric ones and (ii) takes effect size into consideration (Cliff's Delta). Menzies' changes make the SK algorithm suitable to be used in the context of large datasets, where p-values are close to meaningless without an accompanying effect size measure and few or even no assumptions can be made in terms of the shape of the distributions under evaluation.

**Findings. Observation 1) Internal factors provide little explanatory power in the characterization of transaction processing times.** Our random forest models achieve identical mean and median adjusted  $R^2$  scores: 0.16. These scores indicate that internal factors can only explain approximately 16% of the variance in processing times.

**Observation 2) Gas price, median price of pending transactions, and nonce are the top-3 most important features of the model.** Feature importance ranks are summarized in Fig. 5. It is unsurprising that the *transaction's gas price* (`gas_price_gwei`) is the most important



**Fig. 5** The distribution of absolute Shapley values per feature (minutes). Features are grouped according to their importance rank (lower rank, higher importance). Due to the high skewness of the data, we apply a  $\log(x+1)$  transformation to the y-axis

feature, since the incentive that miners receive to process transactions is a function of the gas price of the transactions that they choose to process. Nevertheless, we reiterate that our model achieves a median adjusted  $R^2$  of only 0.16 (Observation 1). That is, despite its inherent importance, the gas price of a transaction explains less than 0.16 of the variability in transaction processing times. We conjecture that the explanation lies in the *relative* value of a given gas price. In a way, gas prices can be interpreted as bids in an auction system. The actual, practical value of a bid will always depend on the bids offered by others. Hence, it is entirely possible that the same gas price of x GWEI could be regarded as *high* in one day (i.e., when other transactions typically have a gas price that is lower than x) and *low* in

some other day (i.e., when the opposite scenario takes place). We explore this conjecture as part of RQ2, when we evaluate the explanatory power of pricing behavior.

The second most important feature is the *median gas price of recent pending transactions from the same transaction issuer* (`med_pend_prices`). In Ethereum, a transaction  $t$  from a given transaction issuer  $I$  can only be processed once all prior transactions from  $I$  have been processed (i.e., those transactions with *nonce* lower than that of  $t$ ). For instance, a transaction can stay in a pending state for a long time when a prior transaction from the same issuer has a very low gas price. Hence, pending transactions penalize the processing time of the current transaction. In fact, if the value of `med_pend_prices` is higher than zero, then we know by construction that there exists at least one pending transaction (since gas price cannot be zero). We assume that our model has learned this rule. Interestingly, the third most important feature is the *transaction's nonce* (`tx_nonce`). Hence, we believe that our model is combining the second and third most important features to estimate the time penalty induced by pending transactions.

There is also a complementary facet to `tx_nonce`, which is bound to the transaction issuer. A transaction with a high nonce indicates that its transaction issuer has submitted many transactions in the past. Experienced transaction issuers are likely to be more careful in setting up their gas prices in comparison to novice issuers (e.g., they might be more mindful about the competitiveness of their chosen gas prices compared to novices) (Oliva and Hassan 2021).

**Observation 3) Contextual features have little importance.** The most important contextual feature is ranked only 7th. Its median Shapley value is 0.03 minutes (1.8 seconds). In other words, in at least 50% of the cases, this features either adds or removes only 1.8 seconds to the final prediction.

### Summary

**How well do blockchain internal factors characterize transaction processing times?**

Internal factors provide little explanatory power in the characterization of transaction processing times. In particular:

- Our random forest models only achieve both mean and median adjusted  $R^2$  scores of 0.16.
- Gas price, median price of recent pending transactions from the same transaction issuer, and nonce are the top-3 most important features of our model.
- Contextual features have little importance.

## 6 RQ2: How Well does Gas Pricing Behavior Characterize Transaction Processing Times?

**Motivation** The results of our previous research question indicates that internal factors of the Ethereum blockchain only explain a small amount of the variability in transaction processing times.

Similar to the stock market, the Ethereum blockchain ecosystem is impacted by a plethora of external factors, including speculation (e.g., NFTs, Elon Musk Tweets), the global geopolitical situation (e.g., wars, the USA-China trade war), cryptocurrency

regulations (e.g., FED interventions), and even unforeseen events (e.g., the COVID-19 pandemic). Many of these factors are largely unpredictable and difficult to be captured in the form of engineered features. Yet, these factors are known to impact the *gas pricing behavior* of a large portion of transaction issuers (Ante 2021; Binder 2022). For instance, a cryptocurrency ban in a country can cause a massive amount of people to sell their ETH and tokens, causing many DApps to send transactions to facilitate those actions, thus increasing the average gas price being paid for all transactions in the platform. Since the practical value of a gas price depends on the gas price of all other transactions, we conjecture that pricing behavior (and changes in thereof) influences the processing times of transactions.

**Approach** The additional features that we engineer in order to capture gas pricing behaviors are described in Table 2 along with their rationale. In summary, we introduce 42 additional features to our internal feature set. We primarily involve aggregations of the number of previously processed transactions in recent blocks and their processing times (historical features), in conjunction with the gas price of these transactions (a behavioral feature). For instance, one of our features indicates whether a given gas price is very low/low/median/high/very high by comparing it to the gas prices of all transactions included in the prior 120 mined blocks. In a way, our model from RQ2 provides a *relative* perspective on gas prices instead of an absolute one (RQ1).

Our approach is similar to that of RQ1. First, we reuse the model validation approach from Section 4 to determine the explanatory power of our new model (adjusted  $R^2$ ). Next, we compare the difference in adjusted  $R^2$  between the models from RQ1 and RQ2. The rationale is to determine the importance of the additional features (pricing behavior). We operationalize the comparison by means of a two-tailed Mann-Whitney test ( $\alpha = 0.05$ ) followed by a Cliff's Delta ( $d$ ) calculation of effect size. We evaluate Cliff's Delta using the following thresholds (Romano et al. 2006): negligible for  $|d| \leq 0.147$ , small for  $0.147 \leq |d| \leq 0.33$ , medium for  $0.33 \leq |d| \leq 0.474$ , and large otherwise. Finally, to understand the importance of individual features, we plot their distribution of absolute Shapley values and compute SK.

To further understand the role of individual features, we use *Partial Dependence Plots* (PDP). A PDP reflects the expected output of a model if we were to intervene and modify exactly one of the features. This is different from SHAP, where the Shapley value of a feature represents the extent to which that feature impacts the prediction of single dataset instance while accounting for interaction effects. We also compute pair-wise feature interactions to understand the interplay between specific pairs of features. We use the `explainer.shap_interaction_values()` function from the same shap package. This computes the Shapely interaction index after computing individual effects, and does so by subtracting the main effect of the features to result in the pure interaction effects (Molnar 2020).

**Findings. Observation 4) Our new model explains more than half of the variance in processing times.** With the new feature set, our random forest models achieve identical mean and median adjusted  $R^2$  scores: 0.53. In RQ1, the mean and median scores were also identical at 0.16. As the numbers indicate, the pricing behavior dimension leads to an absolute increase of 0.37 in the median/mean adjusted  $R^2$ . Indeed, the difference between the adjusted  $R^2$  distributions is statistically significant ( $p$ -value  $\leq 0.05$ ) with a large ( $d = 1$ ) effect size. We find these scores confirm the robust explanatory power of our features, demonstrating the stability of our models.

In particular, they are compatible (and even sometimes outperform) adjusted  $R^2$  scores reported in software engineering studies that built regression models with the purpose of understanding a certain phenomenon (Bird et al. 2011; Jiarpakdee et al. 2021; McIntosh et al. 2016). For instance, (McIntosh et al. 2016) built sophisticated regression models using restricted cubic splines to determine whether there is a relationship between (i) code review coverage and post-release defects, (ii) code review participation and post-release defects, and (iii) reviewer expertise and post-release defects. In total, the authors built 10 models and obtained adjusted  $R^2$  scores that ranged from 0.20 to 0.67 (mean = 0.46, median = 0.41, sd = 0.17). In comparison, the median adjusted  $R^2$  score of our model is 0.53, which sits higher than the median score achieved by their models. Most importantly, our score was computed on a *hold-out test set*, while theirs were computed on the *train set*. As the hold-out test set is unseen by the model during training, it is generally harder to achieve a higher  $R^2$  on compared to the train set.

**Observation 5) Gas price competitiveness is a major aspect in determining processing times.** Figure 6 depicts the distribution of absolute Shapley values associated with each feature. The most important feature is the *median of the percentage of transactions processed per block, in the previous 120 blocks, with a gas price below that of the current transaction* (`med_pct_below_120`). The gas price per se (`gas_price_gwei`) only ranks 8th, with a small median absolute Shapley value of 0.01.


**Observation 6) There is a gas price at which transactions are processed closest to the fastest speed possible (and thus increasing such a price is \*not\* likely to further reduce the processing time.)** The PDP depicted in Fig. 7 leads to several insights. First, the range of the y-axis indicates that changes in `med_pct_below_120` have a considerable influence on the predicted processing times. Second, we note an inverse relationship between `med_pct_below_120` and the predicted processing time up until `med_pct_below_120 = 50%`. That is, let  $p(B)$  be the lowest gas price used by a transaction inside a mined block B. Let  $p(t)$  be the gas price of a transaction  $t$ . Setting  $p(t)$  such that  $p(t) > p(B)$  is true for at least half of the 120 most recently mined blocks already provides a speed boost in transaction processing time. From `med_pct_below_120 = 50%` onwards, however, the processing time does not decrease anymore.

In summary, (i) increasing the gas price beyond a certain value  $x$  does not tend to reduce processing times any further and (ii) the value of  $x$  is dependent on the prices of the transaction in the prior 120 blocks (otherwise the `gas_price_gwei` would be more important than `med_pct_below_120`).

**Observation 7) Information about the median price of pending transactions and nonce remain important factors.** Looking at Fig. 6 again, we note that the feature *median price of pending transactions from the same issuer* (`med_pend_prices`) and *transaction nonce* (`tx_nonce`) are ranked second and fourth. In RQ1, our baseline model suggests that these same two features were ranked second (the same) and third (one rank lower) respectively. In other words, these two features remain important factors.

We plotted the PDP for *median price of pending transactions from the same issuer* (`med_pend_prices`) and observed that the y-axis range is very constrained (range of 0.01 for 90% of the values of `med_pend_prices`). In other words, the predicted processing time changes very little when `med_pend_prices` is changed and everything else kept constant. Nevertheless, such a feature ranks second according to SHAP (Fig. 6). Hence, we conjecture that `med_pend_prices` has strong interactions with some other feature. In fact, in RQ1 we conjectured that the model could be using both `med_pend_prices`





**Fig. 6** The distribution of absolute Shapley values per feature in our model with additional features. Features are grouped according to their importance rank. Due to the high skewness of the data, we apply a  $\log(x+1)$  transformation to the y-axis

and `tx_nonce` to penalize processing time due to pending transactions. To investigate our conjecture, we calculated feature interactions between `med_pend_prices` and every other feature. The results indicate that the strongest interaction is indeed with `tx_nonce`. We then computed all pairwise feature interactions and noted that the feature pair  $\{med\_pend\_prices, tx\_nonce\}$  is the second-strongest feature interaction in our model, only behind  $\{med\_pct\_below\_120, std\_num\_below\_120\}$  (i.e., the first and third most important features).



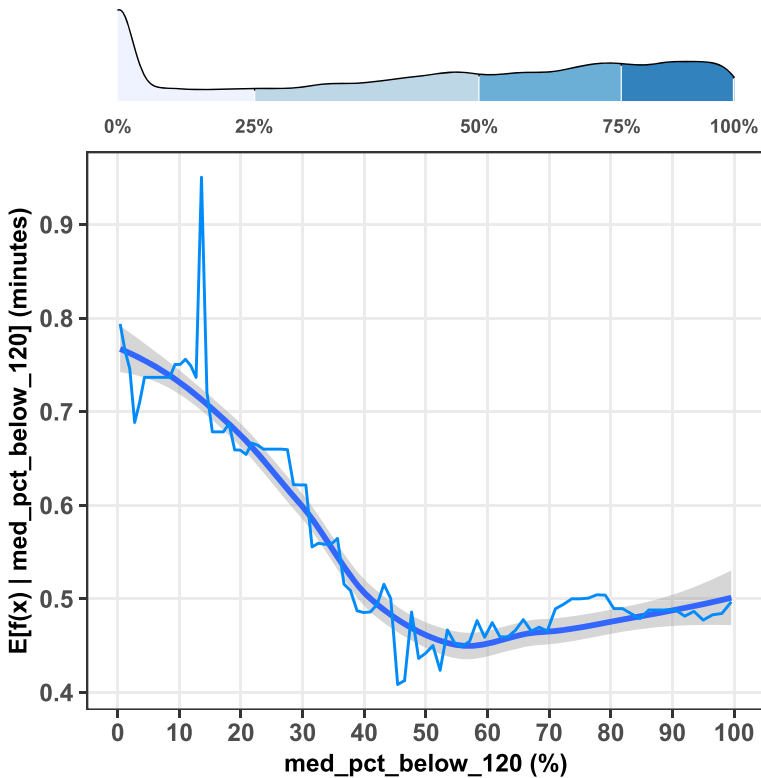


Fig. 7 A partial dependence plot for the med\_pct\_below\_120 feature and transaction processing time

Summary

**RQ2: How well does gas pricing behavior characterize transaction processing times?**

Gas pricing behavior characterize transaction processing times much better than blockchain internal factors alone. In particular:

- Our random forest models achieve a median adjusted  $R^2$  scores of 0.53.
- Gas price competitiveness is a major aspect in determining processing times.
- Setting a gas price above a certain value is unlikely to further reduce processing time.
- Number and gas price of pending transactions have a large influence on the processing time of the current transaction.

## 7 Discussion

In the following section, we discuss the implication drawn from our findings in the previous sections.

**Implication 1)** *ÐApp developers should focus on quantifying how competitive their gas prices are.* Price competitiveness is the most important indicator for transaction pro-

cessing times among our entire set of engineered features. In particular, it is significantly more important than internal blockchain factors that practitioners typically associate with transaction processing times, including: block difficulty (ranks 17 and 20), the number of transactions in the pending pool (rank 10), and network utilization (rank 15). This is interesting considering that it is common for other non-blockchain based transaction processing systems to

In practice, DApp developers should ensure that they execute transactions with competitive gas prices, and should understand the competitiveness of a given specific price is likely to change over time. Given our empirical results, we believe that our metric *med\_pct\_below\_120* can serve as a starting point for DApp developers to quantify price competitiveness. We warn developers, however, that such a price should not be “too high”, since there is a point at which higher prices do not result in lower processing times. More generally, future research in explainable models for transaction processing times should investigate the optimal number of past blocks to consider and how frequently such a number needs to be fine-tuned over time to avoid concept drift.

**Implication 2) *The architecture of DApps should take into account the fact that processing time is only weakly associated with gas usage and gas limit.*** The contribution of work lies in not only identifying those features that are most important but also in identifying those that are not.

Figure 6 indicates that features related to *gas usage*, such as *tx\_gas\_limit* and *std\_func\_gas\_usage\_120* (which is correlated with *med\_func\_gas\_usage\_120*), have little importance. This suggests that transactions with competitive gas prices are prioritized by miners over the amount of gas that they will potentially consume. Hence, if a DApp invokes a heavy computation function that consumes lots of gas, then such a transaction will still require a competitive gas price to be processed in a timely fashion. Such a DApp might be costly and potentially unfeasible to maintain. We briefly discuss two strategies that would help DApp developers achieve a better time-cost balance.

The first strategy is to optimize the source code of smart contracts to make it burn less gas. This problem is known in the literature as gas (consumption) optimization (Zou et al. 2019; Chen et al. 2017; Marchesi et al. 2020). A function that burns less gas would enable transaction issuers to submit transactions using a higher gas price while still maintaining the same overall cost. Consequently, as shown in our results, transactions with more competitive gas prices tend to get processed faster. However, we note that gas consumption optimization should be performed carefully, since such optimizations can hinder program comprehension and testability.

The second strategy refers to mitigating the risk of making a poor gas price choice in the case of transactions that perform repetitive operations (e.g., updating the balance of several user accounts). For instance, assume that the balance of 200k accounts need to be updated. Assume that the DApp in question has a function *f* that takes a list of accounts and updates their balances. Instead of submitting a single, heavy transaction with a potentially poor gas price choice (as current estimation services are far from perfect), it is better to split the workload into batches and send one transaction for each batch. Now, instead of making one single price choice, the transaction issuer needs to make several gas price choices (one for each batch). This amortizes the risk of poor price choices, analogously to the concept of dollar cost averaging as defined by Benjamin Graham in the context of stock investment strategies (Graham and Zweig 2003). Finding opportunities to migrate from a single heavy transaction to several smaller transactions that can be executed in batches often requires the DApp’s architecture to be planned in such a way that contract transactions are as atomic as possible.

We acknowledge, however, that the aforementioned strategies are unfeasible in specific scenarios or use cases. Most importantly, our work highlights an important and yet not so obvious interplay between DApp architecture design and transaction processing times.

**Implication 3)** *We did not observe a strong relationship between contextual features and processing time. However, future research should not discard this type of feature.* We observed that features in the contextual dimensions generally contribute very little to the resulting prediction. Most notably, neither network utilization (*net\_util*) nor network congestion (*pending\_pool*) hold strong relationships with processing time. This is in contrast with conventional transaction processing systems (e.g., cloud-based services), which are likely to have an increase in processing times when many users are simultaneously using the system. This may occur either because (i) contextual features really do not matter that much compared to pricing behavior or (ii) because contextual factors are not strongly present in our dataset (e.g., seasonalities and network congestion). We thus invite future research to further explore contextual features with other datasets and timeframes. Additional contextual features should also be engineered, as they may impact processing times more heavily than those included in our study.

## 8 Related Work

We note that the analysis we conduct in our study on the explanatory power of an exhaustive set of features with respect to transaction processing times has not yet been investigated. We discover several insights using this proposed approach, the majority of which suggest the importance of historical features related to gas prices. We hope our contributions prove useful for DApp developers, and invite future research to build upon our study.

Pierro and Rocha (2019) investigate how several factors related to transaction processing times cause or correlate with variations in gas price predictions made by Ethereum's Gas Price oracle<sup>15</sup>.

The authors conclude there is non-directional causality between oracle gas price predictions and the time to mine blocks, a unidirectional causality (inverse correlation) between oracle predictions and number transactions in the pending pool, and a unidirectional causality (inverse correlation) between oracle predictions and the number of active miners.

In a later study the same authors (Pierro et al. 2022) evaluate the processing time and corresponding gas price predictions made by EthGasStation's Gas Price API. To do this, the authors use processed transactions by retrieving those with a gas price greater or equal to the gas prices in all processing time predictions made by EthGasStation. They then verify whether those transactions were processed in the following  $j$  blocks, as predicted by EthGasStation. The authors conclude that EthGasStation holds a higher margin of error compared to what they claim. Using a Poisson model, the authors achieve more accurate predictions than EthGasStation by considering data within only the previous 4 blocks. In turn, the authors suggest that such models should be constructed using data within the most recent blocks, as opposed to the 100 previous blocks considered by EthGasStation, to improve prediction accuracy.

de Azevedo Sousa et al. (2021) investigate the correlation between transaction fees and transaction pending time in the Ethereum blockchain.

The authors calculate Pearson correlation matrices between several features across 1) the full set of transactions, 2) distinct sets of transactions based on their gas usage and gas

<sup>15</sup><https://etherchain.org/tools/gasPriceOracle>

price, and 3) clusters of transactions resulting from DBSCAN. The authors conclude there is a negligible correlation between all features related to transaction fees and the pending time of transactions.

Singh and Hafid (2020) focus on the prediction of transaction confirmation time within the Ethereum blockchain. The study includes 1 million transactions with non-engineered features from the blockchain. The authors discretize transactions by their confirmation times (15s, 30s, 1m, 2m, 5m, 10m, 15m, and  $\geq 30m$ ). The majority of constructed classifiers (Naive Bayes, Random Forests, and Multi Layer Perceptrons) achieved higher than 90% accuracy, with the Multi Layer Perceptron classifier regularly outperforming the others.

The impact of the transaction fee mechanism (TFM) EIP-1559 on processing times is investigated by Liu et al. (2022) in the Ethereum blockchain since its relatively recent implementation. The authors discover that EIP-1559 makes fee estimation easier for users, ultimately reducing both transaction processing times and gas price paid of transactions within the same block. However, when the price of Ether is volatile, processing times become much longer. We find that features related to historical gas prices are most important in estimating processing time, and considering EIP-1559 makes fee estimations and thus assigning gas prices easier for users, similar TFMs may help DApp developers avoid slow processing times altogether.

The study conducted by Kasahara and Kawahara (2019) uses queuing theory to investigate the impact of transaction fees on transaction processing time within the Bitcoin blockchain. The authors conclude that transactions associated with high transaction fees are processed quicker than those that are not. The authors also find network congestion strongly impacts the processing time of transactions, by increasing the time, regardless if they are associated with either low or high fees, and regardless of block size. Conversely, our study finds network utilization to hold minimal impact on transaction processing time.

Oliva et al. (2020) analyze how smart contracts deployed on the Ethereum blockchain are commonly used. The authors conclude that 0.05% of total smart contracts were the target of 80% of all transactions executed during the time period of their study. In addition, 41.3% of these contracts were found to be deployed for token-related purposes. Our paper includes information regarding the smart contract targeted by a transaction if applicable, though we do not find that these features have strong impacts on transaction processing time.

The work of Zarir et al. (2021) investigates how DApp developers can develop cost effective DApps. Similar to our study, the authors conclude transactions are prioritized based on their gas price, rather than their gas usage. In addition, the authors find the gas usage of a function can be easily predicted. As a result, the authors suggest that DApp developers provide gas usage information in their smart contracts, and platforms such as Etherscan and wallets should provide users with historical gas usage information for smart contract functions.

## 9 Threats to Validity

**Construct validity** The *pending timestamp* of a transaction is not publicly available information within the Ethereum blockchain. As a result, we rely on Etherscan to provide an accurate measurement of when a transaction was actually executed. Because of the size, reputability, and popularity of Etherscan, we continue to rely on it and reaffirm the accuracy of their provided data.

This paper is a first attempt toward using predictive models in order to derive and explain insights from data within the context of the Ethereum blockchain. We choose to leverage

Random Forest models as they hold an optimal balance between performance and explainability. We encourage future work to build upon the approach demonstrated in our study in order to achieve models which achieve higher performance, and thus further contribute to predicting transaction processing times and the generalizability of the model explanations in our study.

Our Random Forest models are hyperparameter tuned in each bootstrap to control states of the model as well as to achieve greater performance on the hold out test set. We only consider a few hyperparameters to control, and there exist multiple hyperparameters that can be tuned which were not considered in our study. We invite future researchers to consider a more exhaustive set of hyperparameters to tune in similar Random Forest models, and to investigate their impact on performance.

Although we opt to leverage the adjusted  $R^2$  metric as a robust approach to determine the goodness of fit of the model, other performance metrics exist that can, in certain contexts, contribute toward a complete assessment of the performance of the model. As such, future work should assess similar models using additional performance metrics.

We fully rely on SHAP values to generate explanations for the predictions of our Random Forest models. It is possible that different feature importance methods would generate different results. We thus encourage future work to explore other methods for explaining models (e.g., LIME (Ribeiro et al. 2016)).

**Internal validity** In our paper, we choose to study the explanatory power of a set of both internal and external features spanning across various dimensions. The majority of features in our set include those which primarily reflect technical aspects of the Ethereum blockchain. However, other types of features, such as those that are based upon and or track market speculation (e.g., Elon Musk tweets (Ante 2021)), exist which might play a big role in describing the behavior of the Ethereum blockchain. As a result, we invite future work to design and investigate an even more exhaustive set of features which were not included in our study. Additionally, we use some features as a proxy for capturing other information, including using the transaction nonce (*tx\_nonce*) as user experience. Instead, future work can explicitly collect such information to study their direct impact on processing times

The collected data and concluding results naturally depend on the defined time frame of our study. We also use sampling methods to construct our results, though we draw a representative sample of blocks appended to the Ethereum blockchain per day, in an attempt to capture the complete and natural behavior of the network during this period. Seasonal trends are known to impact the blockchain (Kaiser 2019; Werner et al. 2020) and therefore may have affected the insights and conclusions derived in our study. As a result, we suggest that future studies should validate our findings by replicating them using different and longer time frames. Finally, we emphasize that in this paper we demonstrate an approach that is extensible, and therefore can be followed in similar contexts with similar goals.

In addition to the dependence on the chosen time frame, our study also relies on actively collecting data from Etherscan. As Etherscan does not offer a robust API which allows us to collect all data points used in our study, we developed other methods of data collection. Our data collection methods respect Etherscan's platform by adhering to standards they have set so as to not disrupt the availability of their normal services. Because of this, we are unable to collect every transaction executed within the Ethereum blockchain, which might result in imbalanced data, such as collecting a low amount of transactions that are processed in extremely fast or slow times.

**External validity** By construction, the conclusions that we draw in our paper are tied to the dataset and time frame that we analyzed. Since our approaches for data curation and

model interpretation were designed to be robust, we consider our conclusions to be reliable. Nonetheless, due to the phenomenon of concept drift (Nishida and Yamauchi 2007), most machine learning models need to be retrained from time to time (a.k.a., periodic retraining, refreshing). In this context, it remains unclear how long it would take for our model to start suffering from concept drift. We invite future research to build models similar to ours and evaluate their resilience to concept drift.

Our study focuses solely on the Ethereum blockchain during a specific time frame. As it is likely for blockchains to be unique in purpose and design, the consistency of the conclusions resulting from our study may not hold in other blockchain environments. For example, transaction fees in the Bitcoin blockchain are calculated differently than fees in the Ethereum blockchain, and thus might not impact processing times as heavily. Additionally, blocks in the Bitcoin blockchain take about 10 minutes on average to be appended to the chain, which might result in factors impacting transaction processing times very little. We again note that we designed the approach in this paper to be as extensible as possible, in order to be reused in similar contexts with similar goals. In turn, we encourage future research to explore ideas and conduct replication experiments similar to ours in other blockchain platforms, as well as different time frames, which may be fruitful.

## 10 Conclusion

DApPs on the blockchain require code to be executed through the use of transactions, which need to be paid for. For DApps to be profitable, developers need to balance paying high amounts of Ether to have their application transactions processed timely, and high-end user experience. Existing processing time estimation services aim to solve this problem, however they offer minimal insight into what features truly impact processing times, as the platforms do not offer any interpretable information regarding their models.

With this as motivation, we collect data from Etherscan, and Google BigQuery to engineer and investigate features which capture internal factors and gas pricing behaviors. We use these features to build interpretable models using a generalizable and extensible model construction approach, in order to then discover what features best characterize transaction processing times in Ethereum, and to what extent. In particular, we discover that metrics regarding the gas price information of transactions processed in the recent past to hold the most explanatory power of all features. In comparison, the studied features which capture gas pricing behaviors hold more explanatory power than any feature dimension found in internal factors alone and combined. Resulting from our findings, DApp developers should primarily focus on monitoring recent gas pricing behaviors when choosing gas prices for the transactions in their applications. Ensuring gas prices chosen by DApp developers are competitive compared to recent gas pricing will ultimately help them achieve a desired level of QoS. In addition, DApp developers should also avoid setting high gas prices that deviate too much from the recent past to avoid diminishing returns. We encourage future work to engineer and investigate other factors that capture gas pricing behaviors, such as the one proposed in this study, to see how they impact model predictions.

As gas price is much more important than gas limit, DApp developers should avoid designing contract functions which consume large amounts of gas. Instead, developers should consider performing gas optimizations to the contracts' source code and adjusting transaction workloads to mitigate poor gas price choices. We invite future work to directly focus on the relationship between smart contract gas consumption, DApp's architecture design, and transaction processing times.

Our results show that blockchain internal factors are only loosely associated with transaction processing times. Nevertheless, there were specific times in history when Ethereum substantially slowed down due to a flooding of transactions (e.g., the CryptoKitties incident (BBC 2017)). Therefore, future work should reinvestigate the factors impacting processing times in the context of seasonalities and heavy workloads.

ÐApp developers can potentially reduce the amount of internal factors being actively monitored, as we found internal factors to generally contribute very little to the resulting prediction. More specifically, internal factors which are we studied in our contextual dimension do not impact the predictions of the models at all. In turn, ÐApp developers can focus more on gas pricing behaviours and a smaller subset of internal factors when aiming to achieve their desired QoS. However, we suggest that future work should not exclude such factors. Future work should focus on these factors explicitly, for example by engineering new contextual features or including those that our analysis did not, which may be more powerful in their impact on processing times.

We also encourage and invite future work to investigate other explainable models for transaction processing times which include additional features and refine the ones we propose to achieve higher performance. Our most robust models achieve an adjusted  $R^2$  of 0.53, meaning almost half the variance is not covered by the many features we collect and engineer.

Finally, we encourage future research to use the results of our experiments to motivate the continuation of empirical studies within the area of blockchain. In relation to our study specifically, possible topics for future research include: i) the construction and interpretation of other types of models which can more accurately predict processing times, ii) similar experiments conducted in blockchain environments other than Ethereum, and iii) a similar or new set of experiments conducted for comparison which use alternative parameters to those in our study, such as time frames, statistical methods and techniques, and additional features.

## 11 Disclaimer

Any opinions, findings, and conclusions, or recommendations expressed in this material are those of the author(s) and do not reflect the views of Huawei.

## Appendix A: Model Construction: Correlation and Redundancy Analysis

**Table 3** The independent features removed resulting from conducting the correlation analysis in Section 4.3

Index	Feature 1 (Removed)	Feature 2
1	num_same_120	avg_num_same_120
2	difficulty_avg_120	difficulty_med_120
3	med_num_above_120	num_above_120
4	num_below_120	avg_num_below_120
5	avg_num_above_120	med_num_above_120
6	avg_num_below_120	med_num_below_120
7	total_txs_120	avg_txs_120

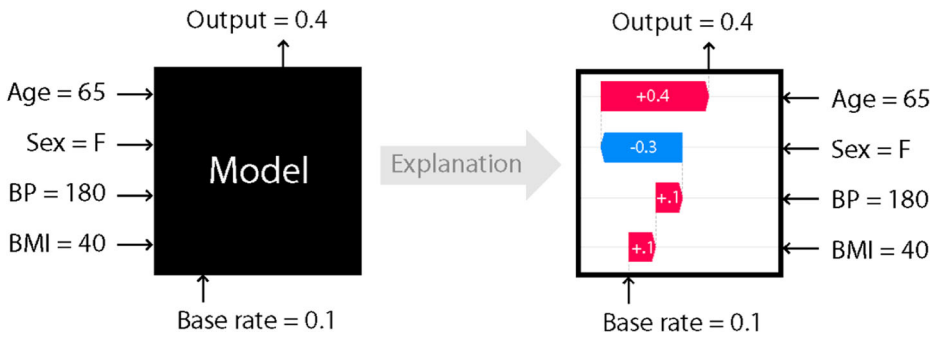
**Table 3** (continued)

Index	Feature 1 (Removed)	Feature 2
8	difficulty_1	difficulty_med_120
9	avg_func_gas_usage_120	med_func_gas_usage_120
10	avg_txs_120	med_txs_120
11	avg_num_same_120	std_num_same_120
12	avg_pending_pool_120	med_pending_pool_120
13	avg_gas_price_1	med_gas_price_1
14	avg_gas_price_120	med_gas_price_120
15	med_num_above_120	med_num_below_120
16	med_gas_price_120	med_gas_price_1
17	pending_pool	med_pending_pool_120
18	med_gas_price_120	gas_price_gwei
19	input_length	to_contract
20	to_contract	contract_bytecode_length
21	avg_gas_price_gwei_prev_day	med_gas_price_120
22	past_avg_mins_120	past_std_mins_120
23	med_func_gas_usage_120	std_func_gas_usage_120
24	contract_bytecode_length	std_func_gas_usage_120
25	num_same_1	pct_same_1
26	pct_above_120	med_pct_above_120
27	avg_pct_above_120	med_pct_above_120
28	pct_below_120	avg_pct_below_120
29	avg_pct_below_120	med_pct_below_120
30	med_num_same_120	med_pct_same_120
31	pct_same_120	avg_pct_same_120
32	med_pct_above_120	med_pct_below_120
33	std_num_same_120	std_pct_same_120
34	med_num_below_120	med_pct_below_120
35	std_pct_above_120	std_pct_below_120
36	avg_pct_same_120	std_pct_same_120
37	num_above_1	pct_above_1
38	num_below_1	pct_below_1
39	std_num_above_120	std_pct_below_120
40	pct_above_1	pct_below_1
41	pct_below_1	med_pct_below_120
42	num_pending	med_pend_prices
43	avg_pend_prices	med_pend_prices
44	std_pend_prices	med_pend_prices

## Appendix B: SHAP

To interpret our random forest models at the feature level, we use a model-agnostic interpretation technique called SHAP (Lundberg and Lee 2017). The intuition behind SHAP





**Fig. 8** Intuition behind the SHAP model interpretation technique (adapted from <https://shap.readthedocs.io>)

is shown in Fig. 8. Assume that we are using a black-box model to predict the chances of someone eventually having a heart attack (HT). The model's features are Age, Sex, BP (blood pressure), and BMI (body mass index). By taking the mean of the HT column, we have an expected value for HT. We refer to this expected value as the *base rate*. In the example shown in Fig. 8 (left-hand side), this base rate is 0.1 (10%). For a particular instance (row) in the dataset (Age = 75, Sex = F, BP = 180, and BMI = 40), the black-box model outputs a HT of 0.4. The question now is: how much does each feature contribute to moving the output from 0.1 (base rate) to 0.4 (the actual prediction)? In an additive model such as a linear model, the effect of each feature is simply the weight of the feature times the feature value. SHAP relies on game theory to extract the features contributions from any machine learning model. As shown in the right-hand side of Fig. 8, SHAP indicates that BMI = 180 contributed with +0.1 to the base rate (positive impact, shown in red), BP = 180 contributed with an additional +0.1 to the base rate, Sex = F contributed with -0.3 to the base rate (negative impact, shown in blue), and Age = 65 contributed the most with +0.4 to the base rate. If we sum up all these individual feature contributions, we obtain the final prediction of HT = 0.4. Hence, SHAP is additive in nature and can easily explain how each feature contributed to the prediction of a given instance in the dataset. SHAP has been increasingly adopted in the Software Engineering community (Viggiato and Bezemer 2020; Esteves et al. 2020; Rajbahadur et al. 2021).

**Data Availability** A replication package containing our studied transaction processing time data is available online in a GitHub repository<sup>16</sup>.

## Declarations

**Conflict of Interests** The authors declared that they have no conflict of interest.

## References

- Aniche M, Maziero E, Durelli R, Durelli V (2020) The effectiveness of supervised machine learning algorithms in predicting software refactoring. *IEEE Transactions on Software Engineering*
- Ante L (2021) How elon musk's twitter activity moves cryptocurrency markets. *Advertising & Marketing Law eJournal*

<sup>16</sup>[https://github.com/SAILResearch/suppmaterial-20-mpacheco-ethereum\\_tx\\_time\\_interpretation](https://github.com/SAILResearch/suppmaterial-20-mpacheco-ethereum_tx_time_interpretation)


- Bao L, Xia X, Lo D, Murphy GC (2021) A large scale study of long-time contributor prediction for github projects. *IEEE Trans Softw Eng* 47(6):1277–1298. <https://doi.org/10.1109/TSE.2019.2918536>
- BBC (2017) Bbc news: Cryptokitties craze slows down transactions on ethereum. <https://www.bbc.com/news/technology-42237162>, [Online; accessed 02-November-2022]
- Bergstra J, Bengio Y (2012) Random search for hyper-parameter optimization. *J Mach Learn Res* 13(2):281–305
- Binder M (2022) Bored ape yacht club caused ethereum fees to soar to astronomical levels. <https://mashable.com/article/ethereum-gas-fees-skyrocket-bored-ape-yacht-club-otherside-nft-launch>, [Online; accessed 10-May-2022]
- Bird C, Nagappan N, Murphy B, Gall H, Devanbu P (2011) Don't touch my code! examining the effects of ownership on software quality
- Boslaugh S, Watters PA (2008) Statistics in a nutshell - a desktop quick reference
- Breiman L (2001) Random forests. *Mach Learn* 45(1):5–32
- Buterin V (2014) Ethereum: A next-generation smart contract and decentralized application platform. <https://github.com/ethereum/wiki/wiki/White-Paper>, [Online; accessed 20-November-2019]
- Chen T, Li X, Luo X, Zhang X (2017) Under-optimized smart contracts devour your money. In: 2017 IEEE 24Th international conference on software analysis, evolution and reengineering, SANER, IEEE, pp 442–446
- Comben C (2018) What are blockchain confirmations and why do they matter? <https://coincentral.com/blockchain-confirmations>, [Online; accessed 04-December-2019]
- de Azevedo Sousa JE, Oliveira V, Valadares J, Dias Gonçalves G, Moraes Villela S, Soares Bernardino H, Borges Vieira A (2021) An analysis of the fees and pending time correlation in ethereum. *International Journal of Network Management* 31(3), <https://doi.org/10.1002/nem.2113>
- Esteves G, Figueiredo E, Veloso A, Viggiano M, Ziviani N (2020) Understanding machine learning software defect predictions. *Autom Softw Eng* 27(3):369–392
- Fan Y, Xia X, Lo D, Hassan AE (2020) Chaff from the wheat: Characterizing and determining valid bug reports. *IEEE Trans Softw Eng* 46(05):495–525. <https://doi.org/10.1109/TSE.2018.2864217>
- Graham B, Zweig J (2003) *The intelligent investor: Revised 1972 Ed.* HarperCollins
- Harrell F (2015) *Regression Modeling Strategies with Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis*, 2nd edn. Springer
- Jakobsson M, Juels A (1999) Proofs of work and bread pudding protocols. In: Proceedings of the IFIP TC6/TC11 Joint Working Conference on Secure Information Networks: Communications and Multimedia Security, Kluwer, B.V., Deventer, The Netherlands, The Netherlands, CMS '99, pp 258–272
- Jiarpakdee J, Tantithamthavorn CK, Grundy J (2021) Practitioners' perceptions of the goals and visual explanations of defect prediction models. In: 2021 IEEE/ACM 18Th international conference on mining software repositories, MSR, IEEE, pp 432–443
- Kaiser L (2019) Seasonality in cryptocurrencies. *Finance Research Letters*, 31
- Kasahara S, Kawahara J (2019) Effect of bitcoin fee on transaction confirmation process. *Journal of Industrial & Management Optimization* 15(1547-5816\_2019\_1\_365):365. <https://doi.org/10.3934/jimo.2018047>
- Kondo M, Oliva GA, Jiang ZMJ, Hassan AE, Mizuno O (2020) Code cloning in smart contracts: a case study on verified contracts from the ethereum blockchain platform. *Empir Softw Eng* 25(6):4617–4675
- Lee D, Rajbahadur GK, Lin D, Sayagh M, Bezemer CP, Hassan AE (2020) An empirical study of the characteristics of popular minecraft mods. *Empir Softw Eng* 25(5):3396–3429
- Liaw A (2010) [r] random forest auc. <https://stat.ethz.ch/pipermail/r-help/2010-October/257208.html>, [Online; accessed 17-November-2021]
- Liu Y, Lu Y, Nayak K, Zhang F, Zhang L, Zhao Y (2022) Empirical analysis of eip-1559: Transaction fees, waiting time, and consensus security. arXiv:220105574
- Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates Inc., Red Hook, NY, USA, NIPS'17, pp 4768–4777
- Lundberg SM, Erion GG, Lee SI (2018) Consistent individualized feature attribution for tree ensembles. arXiv:180203888
- Lyu Y, Rajbahadur GK, Lin D, Chen B, Jiang ZM (2021) Towards a consistent interpretation of aiops models. *ACM Trans Soft Eng Methodol (TOSEM)* 31(1):1–38
- Marchesi L, Marchesi M, Destefanis G, Barabino G, Tigano D (2020) Design patterns for gas optimization in ethereum. In: 2020 IEEE International workshop on blockchain oriented software engineering (IWBOSE), IEEE, pp 9–15
- Mcintosh S, Kamei Y, Adams B, Hassan AE (2016) An empirical study of the impact of modern code review practices on software quality. *Empir Softw Eng* 21(5):2146–2189. <https://doi.org/10.1007/s10664-015-9381-9>

- Menzies T (2020) Scott Knot with nonparametric effect size and significance test. <https://gist.github.com/timm/41b3a8790c1adce26d63c5874fba393>, [Online; accessed 04-May-2021]
- Menzies T, Greenwald J, Frank A (2007) Data mining static code attributes to learn defect predictors. *IEEE Trans Softw Eng* 33(1):2–13. <https://doi.org/10.1109/TSE.2007.256941>
- Midi H, Sarkar S, Rana S (2013) Collinearity diagnostics of binary logistic regression model. *Journal of Interdisciplinary Mathematics* 13:253–267. <https://doi.org/10.1080/09720502.2010.10700699>
- Molnar C (2020) *Interpretable machine learning*. Lulu
- Nishida K, Yamauchi K (2007) Detecting concept drift using statistical testing. In: *International conference on discovery science*, Springer, pp 264–269
- Oliva GA, Hassan AE (2021) The gas triangle and its challenges to the development of blockchain-powered applications. In: *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp 1463–1466
- Oliva GA, Hassan AE, Jiang ZMJ (2020) An exploratory study of smart contracts in the ethereum blockchain platform. *Empir Softw Eng* 25(3):1864–1904
- Oliveira VC, Almeida Valadares J, A Sousa JE, Borges Vieira A, Bernardino HS, Moraes Villela S, Dias Gonçalves G (2021) Analyzing transaction confirmation in ethereum using machine learning techniques. *SIGMETRICS Perform Eval Rev* 48(4):12–15.
- Pacheco M, Oliva GA, Rajbahadur GK, Hassan AE (2022) Is my transaction done yet? an empirical study of transaction processing times in the ethereum blockchain platform. *ACM Transactions on Software Engineering and Methodology*
- Pierro GA, Rocha H (2019) The influence factors on ethereum transaction fees. In: *2019 IEEE/ACM 2Nd international workshop on emerging trends in software engineering for blockchain, WETSEB*, IEEE, pp 24–31
- Pierro GA, Rocha H, Ducasse S, Marchesi M, Tonelli R (2022) A user-oriented model for oracles' gas price prediction. *Futur Gener Comput Syst* 128:142–157
- Rajbahadur GK, Wang S, Ansalidi G, Kamei Y, Hassan AE (2021) The impact of feature importance methods on the interpretation of defect classifiers. *IEEE Transactions on Software Engineering*
- Ribeiro MT, Singh S, Guestrin C (2016) “why should i trust you?”: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, USA, KDD '16, pp 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- Romano J, Kromrey J, Coraggio J, Skowronek J (2006) Appropriate statistics for ordinal level data: Should we really be using t-test and Cohen's d for evaluating group differences on the NSSE and other surveys? In: *Annual meeting of the Florida Association of Institutional Research*, pp 1–3
- Scott AJ, Knott M (1974) A cluster analysis method for grouping means in the analysis of variance. *Biometrics* 30(3):507–512
- Signer C (2018) *Gas cost analysis for ethereum smart contracts*. Master's thesis, ETH Zurich, Department of Computer Science
- Singh HJ, Hafid AS (2020) Prediction of transaction confirmation time in ethereum blockchain using machine learning. In: Prieto J, Das AK, Ferretti S, Pinto A, Corchado JM (eds) *Blockchain and applications*. Springer International Publishing, Cham, pp 126–133
- Tagra A, Zhang H, Rajbahadur GK, Hassan AE (2022) Revisiting reopened bugs in open source software systems. *Empir Softw Eng* 27(4):1–34
- Tantithamthavorn C, Hassan AE (2018) An experience report on defect modelling in practice: Pitfalls and challenges. In: *Proceedings of the 40th International conference on software engineering: Software engineering in practice*, pp 286–295
- Tantithamthavorn C, McIntosh S, Hassan AE, Matsumoto K (2016) An empirical comparison of model validation techniques for defect prediction models. *IEEE Trans Softw Eng* 43(1):1–18
- Tantithamthavorn C, Hassan AE, Matsumoto K (2020) The impact of class rebalancing techniques on the performance and interpretation of defect prediction models. *IEEE Trans Software Eng* 46(11):1200–1219. <https://doi.org/10.1109/TSE.2018.2876537>
- Tantithamthavorn CK, Jiarpakdee J (2021) Explainable ai for software engineering. In: *2021 36Th IEEE/ACM international conference on automated software engineering, ASE*, IEEE, pp. 1–2
- Thongtanunam P, Hassan AE (2020) Review dynamics and their impact on software quality. *IEEE Trans Softw Eng* 47(12):2698–2712
- Viggiato M, Bezemer CP (2020) Trouncing in dota 2: an investigation of blowout matches. In: *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment*, vol 16, pp 294–300
- Werner SM, Pritz PJ, Perez D (2020) Step on the gas? a better approach for recommending the ethereum gas price. In: *Mathematical Research for Blockchain Economy*, Springer, pp 161–177

- Wood G (2019) Ethereum: A secure decentralised generalised transaction ledger byzantium version 7e819ec - 2019-10-20. <https://ethereum.github.io/yellowpaper/paper.pdf>
- Yatish S, Jiarpakdee J, Thongtanunam P, Tantithamthavorn C (2019) Mining software defects: Should we consider affected releases? In: Proceedings of the 41st international conference on software engineering, IEEE Press, ICSE '19, pp 654–665. <https://doi.org/10.1109/ICSE.2019.00075>
- Zarir AA, Oliva GA, Jiang ZM, Hassan AE (2021) Developing cost-effective blockchain-powered applications: a case study of the gas usage of smart contract transactions in the ethereum blockchain platform. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 30(3):1–38
- Zhang H, Wang S, Chen TH, Zou Y, Hassan AE (2019) An empirical study of obsolete answers on stack overflow. *IEEE Trans Softw Eng* 47(4):850–862
- Zou W, Lo D, Kochhar PS, Le XBD, Xia X, Feng Y, Chen Z, Xu B (2019) Smart contract development: Challenges and opportunities. *IEEE Trans Softw Eng* 47(10):2084–2106

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.



**Michael Pacheco** conducts research at the Centre for Software Excellence at Huawei, Canada. Michael received his MSc degree at the School of Computing of Queen's University in Canada under the supervision of Dr. Ahmed Hassan and Dr. Gustavo A. Oliva. In addition to his research focusing on understanding and evaluating transaction processing times of Ethereum, his research interests also include Artificial Intelligence for Software Engineering (AI4SE), and Explainable AI (XAI).



**Gustavo A. Oliva** is an Adjunct Professor at the School of Computing of Queen's University in Canada. Gustavo leads the blockchain research team at the Software Analysis and Intelligence Lab (SAIL). In addition to blockchain, he also conducts empirical studies in the fields of software ecosystems (e.g., npm), software evolution, and explainable AI. Gustavo received his MSc and PhD degrees from the University of São Paulo (USP) in Brazil under the supervision of professor Dr. Marco Aurélio Gerosa. More information at: <https://www.gaoliva.com>.



**Gopi Krishnan Rajbahadur** is a Senior Researcher at the Centre for Software Excellence at Huawei, Canada. He holds a PhD in computer science from Queen's University, Canada. He received his BE in computer Science and Engineering from SKR Engineering college, Anna University, India. His research interests include AI Engineering, Dataset License Compliance Research, Software Engineering for Artificial Intelligence (SE4AI), Artificial Intelligence for Software Engineering (AI4SE), Mining Software Repositories (MSR) and Explainable AI (XAI).



**Ahmed E. Hassan** is an IEEE Fellow, an ACM SIGSOFT Influential Educator, an NSERC Steacie Fellow, the Canada Research Chair (CRC) in Software Analytics, and the NSERC/BlackBerry Software Engineering Chair at the School of Computing at Queen's University, Canada. His research interests include mining software repositories, empirical software engineering, load testing, and log mining. He received a PhD in Computer Science from the University of Waterloo. He spearheaded the creation of the Mining Software Repositories (MSR) conference and its research community. He also serves/d on the editorial boards of IEEE Transactions on Software Engineering, Springer Journal of Empirical Software Engineering, and PeerJ Computer Science. More information at: <http://sail.cs.queensu.ca>.