

Analytical Engines With Context-Rich Processing: Towards Efficient Next-Generation Analytics

Viktor Sanca

EPFL

Lausanne, Switzerland

viktor.sanca@epfl.ch

Anastasia Ailamaki

EPFL

Lausanne, Switzerland

anastasia.ailamaki@epfl.ch

Abstract—As modern data pipelines continue to collect, produce, and store a variety of data formats, extracting and combining value from traditional and context-rich sources such as strings, text, video, audio, and logs becomes a manual process where such formats are unsuitable for RDBMS. To tap into the dark data, domain experts analyze and extract insights and integrate them into the data repositories. This process can involve out-of-DBMS, ad-hoc analysis, and processing resulting in ETL, engineering effort, and suboptimal performance. While AI systems based on ML models can automate the analysis process, they often further generate context-rich answers. Using multiple sources of truth, for either training the models or in the form of knowledge bases, further exacerbates the problem of consolidating the data of interest.

We envision an analytical engine co-optimized with components that enable context-rich analysis. Firstly, as the data from different sources or resulting from model answering cannot be cleaned ahead of time, we propose using online data integration via model-assisted similarity operations. Secondly, we aim for a holistic pipeline cost- and rule-based optimization across relational and model-based operators. Thirdly, with increasingly heterogeneous hardware and equally heterogeneous workloads ranging from traditional relational analytics to generative model inference, we envision a system that just-in-time adapts to the complex analytical query requirements. To solve increasingly complex analytical problems, ML offers attractive solutions that must be combined with traditional analytical processing and benefit from decades of database community research to achieve scalability and performance effortless for the end user.

Index Terms—analytics, machine learning, data integration, query optimization, hardware-conscious processing

I. INTRODUCTION

One of the main goals of data analytics is to extract insights, value, and intelligence from the available data sources. Traditionally, in the case of relational database management systems, the data would be stored in tables over which users declaratively specify what they would like to extract from the data, with several layers of abstractions performing various logical and physical optimizations. This takes the complexity, and the responsibility from the hands of the user and relieves them from specifying how the data processing pipeline should execute.

However, we are witnessing and participating first-hand in producing context-rich data, such as documents, text, images, and videos, through social media or other means. This trend is growing and expected to continue with the adoption of metaverse-like technologies [1]. The underlying context, such

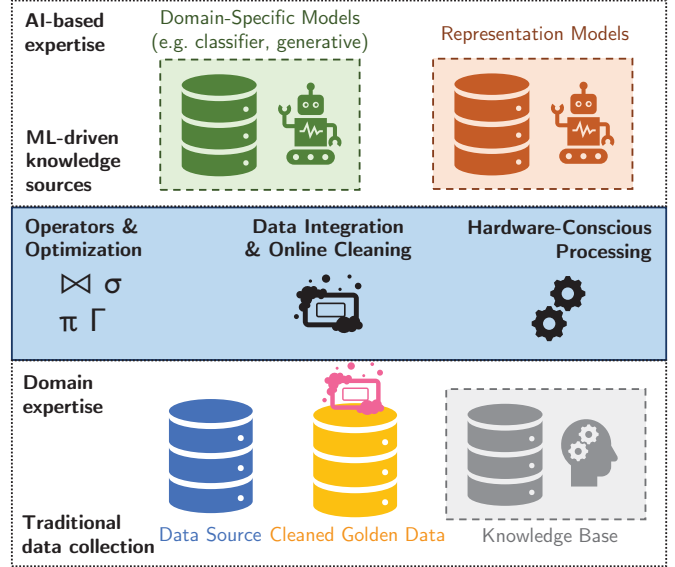


Fig. 1. Analytics with traditional and model-driven sources and processing.

as sentiment or interest detection from text or object recognition and classification from image sources, represents a rich source of value, however unreachable and therefore called *dark* to the data processing tools.

While traditional database systems lack the means to analyze such unstructured data, the explosion of research in machine learning has provided both algorithms and systems fueled by large amounts of data and the latest technology to enable super-human performance in various inference and classification tasks. Tasks previously requiring domain experts and significant intervention of the human-in-the-loop can take place automatically using artificial intelligence. Also, models such as GPT-3 [2] can also represent data sources, generating new data. More broadly, Foundation Models [3] provide an opportunity to adapt large-scale models to specific tasks. Overall, machine learning provides us with analytical capabilities to process and transform the data with human-like context-rich semantics while performing at super-human performance, concerning the speed of processing and the quality of the result.

We group data and insight sources into two categories

in Figure 1. The upper part depicts the sources whose analysis is ML-driven and often have their software systems run individually. Similarly, we consider a traditional data source structured data, collected from various sources (sensors, transactions, domain experts) and cleaned to a certain extent. In a complex data processing pipeline that needs to combine two sources, the user is faced with complex decisions, departing from declaratively specifying only *what* they would like to achieve. Such manual processing is tedious, error-prone, or inefficient and would require a high level of system- and programming- in addition to domain-expertise. To further exacerbate the issue, manually tuning the systems for the particular query to run on increasingly heterogeneous hardware is difficult and impractical.

The challenge for performing optimizations and combining data from various sources (either collected or trained on) is that the data cannot be fully cleaned and unified to an expected schema and values ahead of time. Furthermore, the source of dirty data is less likely to be a mistake such as misspelling but a word with the same semantics (synonym, alternative spelling, alternative forms) - further representing a context-rich online data cleaning task.

To outline the design goals and integration challenges, opportunities, and potential solutions to integrate context-rich processing with analytical engines (the blue-shaded part of Figure 1), in the rest of the paper, we describe:

- a motivating example in section II,
- methods for processing context-rich data in section III,
- integration of the context-rich data sources expert-free and on-the-fly in section IV,
- exposing logical and physical optimizations in section V,
- the challenge of optimizing for hardware heterogeneity in section VI.

We envision an analytical engine that declaratively combines context-rich processing with traditional data sources to hide from the user the complexity of logical and physical optimization, underlying hardware, and resulting on-the-fly data integration.

II. MOTIVATION

Suppose an online shopping platform wants to analyze the data and trends using several sources:

- 1) traditional RDBMS where product, user, and transaction details are managed and stored,
- 2) a general knowledge base to supplement and extend the product information based on domain expertise,
- 3) image storage of the products (from reviews, other websites, or social media).

The data analyst wants to answer which *clothing* products with a price greater than 20 appear in customer images taken after a specific *date*, but such that it is not only the picture of the clothing apparel but other objects appear too - possibly indicating that it is not a generic product description image. This query is conceptually represented in Figure 2.

There are three distinct tasks that a data analyst has to perform:

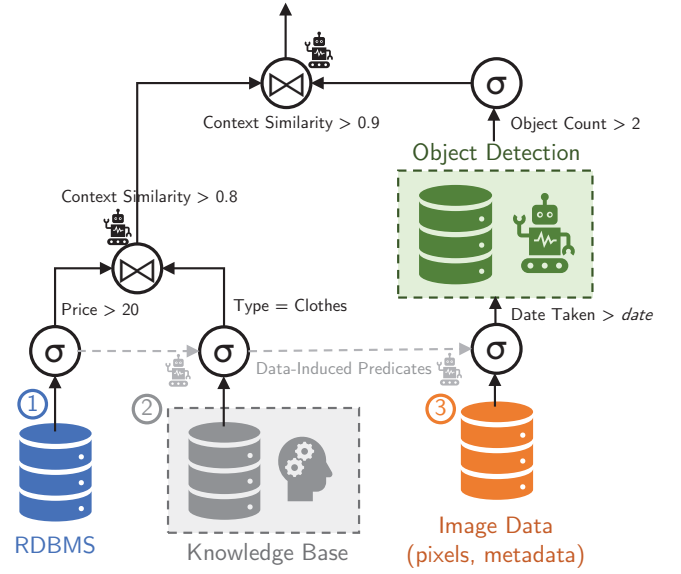


Fig. 2. A query that combines both context-rich and traditional analytics.

- 1) extract the qualifying data from RDBMS,
- 2) infer which ones are clothing items and combine the results with matching clothing objects in images,
- 3) analyze the photos taken after a specific date if they contain more than two objects.

With an RDBMS, step 1) is entirely declarative, and the user can retrieve and export the qualifying records efficiently. However, if not careful, the user may not push down the filter in the knowledge base to reduce the qualifying records for the first join, where only clothing items should be selected. Furthermore, the knowledge base may have been curated and collected on a different and broader dataset that does not precisely match the labels or available data in the RDBMS. These differences are more likely to occur due to synonyms or some other alternative word or phrase forms that may not be resolvable with traditional string similarity algorithms and where the user may need to manually or with the help of an expert resolve which items are indeed clothing items, which can be both error-prone and heavily involve human in the loop which increases the effort and execution time. Finally, the analyst may make a mistake, and instead of first filtering out images after a specific date, perform heavy processing on all the corpora and similarly have issues merging the object labels with ones coming from steps 1) and 2).

Overall, the analyst now has the task to carefully orchestrate possibly an ad-hoc query and multiple sources of data and system components. The cost of logical or physical optimization mistakes, as well as not considering scalability and potentially not processing data beyond main-memory capacity, can quickly hinder even a reasonably simple context-rich query and make it frustrating to use and design such pipelines, abandoning the potential value of the data. Furthermore, every query becomes a hand-tuned orchestration between the

components, including logical optimization steps, physical optimization, and hardware resource management, despite all the lessons learned from decades of data management research.

Several requirements surface from the example, reiterating the need for declarative processing where users do not need to specify *how*, but only *what* they want to process. To have a declarative context-rich analysis system, we propose a design that necessitates the following:

- 1) cost- and rule-based optimization,
- 2) hands-free hardware resource management,
- 3) ad-hoc data matching and integration,
- 4) exposing automated context-rich processing methods.

The broad database and machine learning research community have focused actively on solving some or all of these issues. In the next section we will discuss the need for automated context-rich data analysis methods.

III. ANALYZING CONTEXT-RICH DATA

We broadly use the term context-rich data for all the data sources that would require a human or a domain expert to infer the meaning (context) that is not immediately stated as machine-understandable rules. For example, string edit distance or locality-sensitive hashing-based string similarity can compare strictly specified characteristics, but such methods cannot capture string synonyms in a straightforward way, where one would use a dictionary or a human-in-the-loop. Similar can apply to images, where finding the most prevalent colors is easy to specify and compute; however, understanding the image setting and performing object detection and classification would be considered context-rich analysis.

One way to perform the such large-scale analysis is to deploy crowdsourcing approaches such as Amazon Mechanical Turk [4] or reCAPTCHA [5], or employ domain experts. These approaches have proven invaluable in annotating data and create sources that were, for example, immensely useful in training machine learning models. However, the human-based analysis would be error-prone, slow, and expensive for analyses.

On the other hand, machine learning had its renaissance due to the data volume explosion, hardware processing power improvement, and novel systems and algorithms, which enabled creating of tools that outperform humans even at context-rich processing tasks. Tasks such as text, image, audio, and video analysis are ongoing topics of research, and specific tasks yield even super-human performance. More importantly, obtaining high-quality models trained on Web-scale data as a commodity resource that can be modified via mechanisms such as one- or few-shot learning is possible, for example, via the idea and initiative of Foundation Models [3]. Context-rich processing must happen seamlessly and scale with the data and request volume, and model-based approaches allow high-recall analysis on various tasks.

It is desirable to use models pre-trained on web-scale data with significant compute resources and fine-tune them to the particular task or dataset at hand. Creating and using models and running them efficiently is possible through frameworks

such as Tensorflow [6] or PyTorch [7], with further extensions such as Keras [8] to further make models more declarative.

There has been a significant interest in bringing together statistical processing with ML training and inference into the traditional data management systems, using systems such as Raven [9], MADLib [10], or SystemML [11] to perform inference or learning on tasks such as regression, classification, clustering to extract richer insights from existing, often tabular data. The lessons and the need are the same: we need viable integration for context-rich analysis and logical and physical optimizations for new data sources and value to avoid the pitfalls of manually constructing complex unoptimized data processing pipelines.

Since the models have been trained on large-scale data, and conversely, we can use data to enrich the analytical pipeline from various sources such as knowledge bases, we cannot assume that data is integrated and cleaned even in the outputs. Cleaning the data ahead of time is impractical, if not impossible, due to the scale. Furthermore, models are probabilistic and do not produce only labels (which can be context-rich on their own), for example, as in Table I, but generative models can produce output and data on their own [2]. Integrating the data between the analytical pipelines becomes a key challenge since online result consolidation typically remains the only option [12].

TABLE I
EXAMPLE OF CONTEXT-RICH TEXT LABELS THAT MODELS MAY OUTPUT

category	semantic matches
dog	dog, canine, golden retriever, puppy
cat	cat, maine coon, feline, kitten
animal	cat, dog, golden retriever, feline
shoes	boots, sneakers, oxfords, lace-ups
jacket	blazer, coat, parka, windbreaker
clothes	boots, parka, windbreaker, coat

However, learned approaches allow us to observe this as a representational learning problem that would allow automated resolution, grouping, joining, and filtering based on context-similarity (such as synonyms and similar categories). Approaches such as word2vec and [13], fastText [14] create a latent representation of strings to learn word associations from large corpora of text, such as Wikipedia [15] or Common Crawl [16]. Such approaches can match similarities based on synonyms, semantics, alternative spellings, different tenses, and even misspellings [17]. Context-based filtering, joining, and grouping based on similarity is then converted to finding similarity between original data representation in the latent, high-dimensional vector space - a task that can be specified and automated.

In the next section, we will discuss how using models to analyze the data and combining the output from different parts of the query can be expressed as operators amenable to logical and physical optimizations across different data processing pipelines.




Fig. 3. Automated, on-the-fly result consolidation.

IV. QUERY PROCESSING AND DATA INTEGRATION

To enable a variety of logical and physical optimizations of the holistic query pipeline, no matter where standard relational or model-based operators are placed, we need to start from a common intermediate representation (IR). Approaches such as Weld [18] or Raven [9] already take significant steps towards a common representation that is amenable to optimization rules.

Since we are to use external models in the pipeline, we need to express some properties of context-rich analysis operators. We need to specify the expected input and output schema at the logical level. Furthermore, we need to select the transformations such operator applies on the input, in terms equivalent to the selection, projection, union, or joins, and what effect operator push-down/pull-up has on other, for example, relational operators. Finally, in this step, we must include high-level cost information, such as the effect on the input/output cardinality. This ensures semantic equivalence and correctness and can be expressed in existing query optimization frameworks [19].

Next, we consider some physical optimizations related to context-rich operators using models. While some models perform computationally-intensive inference only, there are still models that perform similarity search or other methods which impose different costs based on data access patterns that we need to enumerate. For example, index-based access for similarity search [20] should be accounted for in the optimization process in a cost-based optimization process.

Finally, we must process and optimize by considering the potential polystore situation where we will have multiple sources of data and multiple processing engines. To add to this complexity, we also have increasingly heterogeneous hardware at our disposal (CPU, GPU, TPU, FPGA, ASICs, DPU, etc.), all with corresponding data transfer, processing, and startup costs. We need to annotate the processing unit or operator traits and perform decisions that adapt to the underlying hardware, with principles described in the thesis of Chrysogelos on efficient analytical processing for CPU-GPU platforms [21]. Just-in-time decisions are needed to make correct decisions in growing hardware, operator, and system heterogeneity. To fully use the underlying hardware capabilities, this merits optimizing novel analytical operators individually for existing or new platforms and holistically for hybrid execution over a variety of compute units.

Irreplaceable components of context-rich analysis are

context-rich *semantic operators* that can process, transform, and integrate data with traditional analytical operators. This task resembles data cleaning but has to capture context and semantics, perform the task efficiently at query time, and run without user supervision in an automated way. In particular, consider the example in Table I: if we want to join clothing items, without strong domain expertise or knowledge-base, we would not be able to perform matching, and still, data sources may contain synonyms, spelling differences, or other expert-system-detectable differences. To resolve this issue, we propose model-assisted context-rich processing that we call *semantic join*, which we denote with a small robot in Figure 2.

In practice, we can use representational models such as fastText [14] to capture the context even in place of spelling mistakes [17]. More complex models have the capacity to specialize in performing this task with more complex input, such as sentences [2], [22], and work with more complex data, such as images or audio. In particular, model-assisted operators transform the input to latent, high-dimensional vector space, where distance-based metrics such as cosine similarity capture features such as context similarity [13]. We represent this in Figure 3, where a tedious and domain-expert task becomes completely automated, allowing on-the-fly result consolidation based on context.

Traditional techniques such as locality-sensitive hashing or other distance metrics capture syntactic similarity (such as misspellings); however, previously mentioned deep-learning methods learn the context from large data corpora. However, other optimizations known to speed up similarity operations, such as joins, are still applicable and should be explored in this context, where initial embeddings are already provided instead of created based on transformations such as hashing.

We propose the following operator extensions:

- **Semantic Select** - context-based filtering, e.g.
word = "Clothes" using model "M" with
cosine threshold ≥ 0.9 ,
- **Semantic Join** - joining two or more relations based on join key context, for example based on model latent space distance between the join key mappings,
- **Semantic GroupBy** - on-the-fly clustering of the result based on a model-based similarity threshold.

With semantic operators, more complex optimization techniques that work for relational data, such as data-induced predicates, can be evaluated and applied in the query plans [23]. The key idea is to make the specified unified query plan amenable to traditional relational plan optimizations.

We note that SQL may not be the best or the only way to represent such query plans; however having an optimizable representation is required to systematically perform a variety of transformations that preserve the correctness of the execution. Systems such as Weld [18], or effort to push Pandas down the DBMS [24] as a common data science tool - should be considered in such a next-generation system to keep it familiar and interoperable with existing libraries and code-ecosystem.




Fig. 4. Additive effects of logical and physical optimizations (log-scale).

V. QUERY OPTIMIZATION

Cost- and rule-based optimizations must be considered to produce a scalable, efficient, and hardware-conscious data-analytics pipeline. Furthermore, selecting the correct physical operator implementation and scaling the solution to available compute resources improves hardware utilization and end-to-end execution time.

We briefly analyze a model-assisted semantic similarity join to demonstrate the importance of optimizations. To simulate the operation, we join two arrays of 10k strings taken randomly from the Wikipedia dataset. We use fastText word embeddings with a dimension of 100, perform the similarity join with cosine distance where the threshold has to be greater than 0.9, and run the experiment on a server with two socket Intel Xeon CPU Gold 5118 CPU (2 x 12-core, 48 threads) and 384GB of RAM. We summarize the execution time and the effects of optimizations in Figure 4.

Suppose the data analyst takes the first tool at their disposal and uses Python: they load the data eagerly (that may not fit in memory ordinarily) and then specify the task using the available libraries to iterate over two loops and perform a similarity check. Suppose this is a part of a more complex pipeline, and the user did not apply the filter before processing - a standard filter pushdown rule. This represents the left-most bar, and the variant with 1% selectivity applied beforehand in the blue patterned bar. This optimization already improves orders of magnitude even in Python in a single-threaded execution context. Since fastText produces a hash table of known words, we can further try to optimize the amount of data access by prefetching necessary data, a physical optimization detail that the user may not be aware of. The next step takes the execution to tighter code, and fewer function calls between libraries using C++, tightening the similarity code using CPU-specific instructions, and scaling up the execution. The effect of well-known optimizations in the database community brings many orders of magnitude improvements from the initial manual, intuitive solution. Conversely, applying such optimizations

is tedious and error-prone if done manually; however, the frustration when waiting for an answer that takes thousands of seconds which can take less than a second may cause the user to abandon such analysis.

Inefficiencies can quickly add up in a more complex plan, as does the complexity of manually optimizing and tuning ad-hoc context-rich queries. Intuitively, performing expensive model inference or processing benefits equally, if not more, from correct join orders and filter pushdowns that potentially reduce the input cardinality.

Exposing all the operators to optimizations and how they affect cardinality and the input/output characteristics is a necessary prerequisite for the optimization. Frameworks such as Apache Calcite [19] already support flexible operator characteristics, traits, and transformations. We could encapsulate such operators in a UDF-like manner while exposing details such as compute requirements, amenability to parallelizing the input, and memory and data transfer requirements to the optimizer component. Often, there are already highly optimized runtimes to perform inference [6] that can reduce development time and be fused more tightly in global processing orchestration. However, we have to be mindful that some models that are useful for context-rich analysis, such as word embedding models, have both inference modules and data structures such as hash tables and index structures [20] for expediting operations such as similarity or top-k searches, which have to be included in the optimization process equally as relational data indexes are.

The key to optimizing a complex context-rich query, such as one in Figure 2, is to provide correct abstractions that perform the optimizations while enabling the user to declaratively and in an intuitive way specify the data processing task at hand, possibly even making well-educated or learned decisions on some of the semantic operators without explicit user control.

VI. HARDWARE-CONSCIOUS OPTIMIZATIONS

Machine learning workloads were one of the primary drivers of mainstream GPU adoption, which was consequently explored in the context of data management on heterogeneous hardware. Increasingly heterogeneous hardware, with various computational, data movement, and storage access overheads, is a source of active research. The general problem becomes how to efficiently scale up the individual algorithms, tune them to hardware capabilities, orchestrate hybrid execution models, and schedule many machine-learning inference runtimes on the available hardware.

We represent the conceptual heterogeneous hardware layout in Figure 5: it includes multi-core CPUs on multiple sockets, potentially each with a separately interconnected GPU. Furthermore, a TPU-like device is designed to perform fast and efficient inference [25], however, that can equally be considered as an accelerator for relational query processing [26]. It is worth noting that specialized hardware units for processing real-time neural network inference are becoming commonplace even in mobile phones, such as Apple Neural Engine or similar [27]. Finally, there are fast network (e.g.,




Fig. 5. The increasing hardware heterogeneity is challenging to optimize for.

InfiniBand) and storage (NVMe) solutions, all interconnected with PCIe or other technologies. The big challenge is how to provision these resources correctly, how to place, split, and schedule the execution such that the increasingly complex processing pipeline efficiently uses the available resources, and for example, which lessons from hybrid CPU-GPU processing [21] we can take and further generalize.

With increasingly difficult cost and cardinality estimation, fast sampling running on modern hardware [28] or speculation techniques [29] can come in handy to provide mechanisms for practical and adaptive query optimization and execution.

Late binding to the query requirements concerning hardware and data layout requirements has become a standard in modern analytics. Just-in-time code generation using frameworks such as LLVM enables specializing the code paths and equally adapting to an increasing variety of data input and output formats. This is increasingly important, as we now consider not only typical relational formats but repositories that contain context-rich data such as strings, text, audio, images, and video, where the NoDB paradigm [30] and adapting to data format heterogeneity at runtime [31] play a crucial role in handling lazy data access to a variety of data formats. Just-in-time code generation allows this to be specified as late as query runtime: for example, if model operator output specification is not known ahead of time, this can be detected, and the code path generated after the model outputs first data - with corresponding overheads and delay.

Data layouts, data size, compression effects, and access method selection become a factor to consider. For example, recent research has yielded specialized databases for vector processing [32]. Integrating complex pipelines through a well-coordinated polystore mechanism or a fully generated pipeline should take different data formats as a first-class problem. For example, the effects of cache hierarchy or hardware-specific vectorized instructions can significantly change the behavior of processing a high-dimensional vector. Furthermore, optimization opportunities such as inference using hardware-enabled half-precision (or lower) floating point formats need to be considered and specialized for [33].

Finally, data movement and compute shipping have to be carefully designed. Obvious questions such as data compression before sending the data over the interconnect for processing come to mind. However, we should also consider the cost of shipping and initializing model parameters and runtime on devices as a potential cost and optimization, as complex models can have many millions of parameters [3]. The decision of the concrete operator execution happens in a managed runtime such as Tensorflow [6] or being optimized entirely using a relational pipeline [9], further opening up the landscape of scheduling and optimizing complex data-science pipelines to use available heterogeneous hardware resources adaptively based on query-specific requirements.

VII. RELATED WORK

The tasks analytical engines are supposed to support and optimize are evolving with novel algorithms, sources of values, applications, and hardware platforms. The research community and industry are proposing systems and investing significant research efforts to address them. We can use them as foundation blocks for next-generation analytical engines for context-rich processing.

A. Querying and Optimizing Complex Analytical Pipelines

To extract deeper insights from the data, more complex data transformations using User Defined Functions (UDFs) or external libraries can take place. In our case, context-rich analysis can happen as a UDF or by invoking another framework optimized to perform specific tasks (e.g., image object detection). A mounting challenge is optimizing the external operators in the query, where systems like Froid [34] optimize UDFs for Relational DBMS, Weld [18] reduces the cost of data movement through optimizing a standard IR, Raven [9] optimizes ML inference with an RDBMS, while Caesar [35] aims to minimize the overheads between application boundaries. MADLib [10] or SystemML [11] allow training and inference over the input, often tabular data. Frameworks such as TensorFlow perform dataflow optimizations for model training and inference [36]. We envision a system not limited to using models representing only the data in RDBMS but can cross-optimize and use the models from different data sources in a polystore-like design. Our vision aims to combine insights and value from, for example, relational data, knowledge bases, and images in a single analytical pipeline and holistically optimize the analytical pipeline across models and frameworks.

Rather than specifying the queries using languages such as SQL, using a standard Python data analysis library like Pandas and the work on how to expose it to RDBMS optimizations [24] could be a good step to improve usability and reduce the potential complexity of specifying new operators that would be needed for context-rich analytical engine. The end goal would be taking the existing code in languages such as Python and optimizing it automatically by creating a standard IR (intermediate representation), detecting physical and logical optimizations, and correctly tuning the operators to the

underlying hardware capabilities, computation optimizations, and data transfers.

B. System Heterogeneity and Data Integration

The analytical engine components that can perform context-rich analysis are part of reality and exist in individual components. Our vision is to bring them together in an efficient engine that would relieve users from performing tasks other than declaratively specifying what they would like to analyze in the data. This is along the proposal of Dataspaces [37], [38] which proposed a data integration platform. We enhance this fundamental system desiderata and propose how modern hardware, programming language, and machine learning advancements allow novel system interactions and future directions for the intersection of these topics with data management research.

Foundation Models [3] offer a way to democratize models that would be practically impossible to retrain individually due to the required data and computation scale. However, such models can be specialized for various tasks to perform better than a human counterpart on context-rich data. Representational Models such as word2Vec [13] and fastText [14] allow learning the word context and reasoning about such abstract similarity by embedding the text and performing similarity operations in latent space.

Since some of the models are already trained on web-scale data, and the analyst will likely supplement the knowledge with external data sources such as knowledge bases, performing data cleaning tasks on all the input data is impractical, if not impossible. However, data cleaning remains crucial in integrating data from different sources. To this extent, we need to incorporate relevant data in an online rather than ahead-of-query offline process through a query-driven mechanism previously explored in the context of repairing functional dependency violations [12].

C. Adapting to Data and Hardware Heterogeneity

New applications, including machine learning, drive an increasingly heterogeneous hardware landscape. From the CPU-dominant landscape, we have a variety of compute units such as CPU, GPU, TPU, DPU [39], FPGA, ASICs, and others - not only on high-end servers but also on commodity devices such as mobile phones, such as Apple Neural Engine [27]. This is supplemented with various storage devices and interconnects that add to the opportunities for hybrid execution on multiple hardware components and novel optimizations for out-of-core algorithms.

To adapt to the hardware heterogeneity and complex requirements, analytical engines such as Hyper [40], NoisePage [41], and Proteus [42] deploy just-in-time decisions via code generation using frameworks such as LLVM [43]. On the side of machine learning, systems such as TensorFlow use JIT-compilation and specialize the runtime graph utilizing an approach called XLA using a domain-specific compiler.

There has been a significant effort in developing optimization frameworks, such as Apache Calcite [19] for performing logical, rule- and cost-based optimization over heterogeneous

data sources. At a higher level of abstraction, DataCalculator [44] aims to automate the data structure design to hardware and access characteristics, and is a step forward to instance-based optimization without manual tuning, desirable in complex system interactions. Automating the end-to-end system optimization and adapting to data and workload has recently been explored in analytical engines in SageDB [45]. Furthermore, at a lower level of optimization, frameworks such as MLIR [46] provide a hybrid IR (intermediate representation) which can aid in bringing together different system and software components and optimizing them for heterogeneous hardware under unified infrastructure. Holistic system adaptation remains challenging due to data, hardware, and system heterogeneity, creating a vast optimization and interaction space that fuel future and ongoing research directions.

VIII. TOWARDS NEXT-GENERATION ANALYTICS

In addition to the typical structured data analytics, humans produce and collect data that can contain value locked behind rich context for a traditional analytical engine. This may further get exacerbated by proposed use cases such as metaverse, where combining analytical insights from sources of interaction such as images, text, audio, and traditional data sources presents an untapped source of insight.

While context-rich analysis, such as object detection for images or sentiment analysis in text, has proposed solutions using machine learning techniques, these solutions are often optimized individually. For more complex analytical pipelines, the user would perform physical, logical, and hardware optimizations for potentially ad-hoc analytical queries, which we want to avoid through a unified, next-generation analytical engine that applies the lessons learned in data management.

With many specialized and siloed solutions, both ML and DB communities make strides to enable richer and more efficient solutions to extract insights and value from data. We propose a system that performs holistic data processing pipeline optimizations merges the potentially unexpected or dirty results on the fly, and relieves the user from making decisions over the increasingly complex system and hardware landscape.

We envision a next-generation analytical engine that generalizes analytical operations driven by new data processing methods to enable novel ways to extract and combine insights from different sources. At the same time, users are given control of data integration through a new class of semantic and model-assisted relational operators that enable automated result consolidation. Through vertical optimizations, correct abstractions, and just-in-time adaptation, a single declarative framework aims to reach the performance of individually and manually tuned components.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their thoughtful comments. We also thank Panagiotis Sioulas, Anna Herlihy, Eleni Zaprudou, and our friends and colleagues from DIAS lab for their valuable feedback and constructive discussions.

REFERENCES

- [1] E. Burgener and J. Rydning, “High data growth and modern applications drive new storage requirements in digitally transformed enterprises,” *IDC Report: A White Paper, sponsored by Dell Technologies and NVIDIA*, no. IDC doc. #US49359722, Jul 2022. [Online]. Available: <https://www.delltechnologies.com/asset/en-us/products/storage/industry-market/h19267-wp-idc-storage-reqs-digital-enterprise.pdf>
- [2] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020.
- [3] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. S. Chatterji, A. S. Chen, K. Creel, J. Q. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajah, L. Fei-Fei, C. Finn, T. Gale, L. Gillespie, K. Goel, N. D. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. W. Koh, M. S. Krass, R. Krishna, R. Kudipudi, and et al., “On the opportunities and risks of foundation models,” *CoRR*, vol. abs/2108.07258, 2021. [Online]. Available: <https://arxiv.org/abs/2108.07258>
- [4] “Amazon mechanical turk.” [Online]. Available: <https://www.mturk.com/>
- [5] L. Von Ahn, B. Maurer, C. McMillen, D. Abraham, and M. Blum, “reCAPTCHA: Human-based character recognition via web security measures,” *Science*, vol. 321, no. 5895, pp. 1465–1468, 2008.
- [6] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [7] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [8] F. Chollet et al., “Keras,” <https://keras.io>, 2015.
- [9] K. Karanasos, M. Interlandi, F. Psallidas, R. Sen, K. Park, I. Popivanov, D. Xin, S. Nakandala, S. Krishnan, M. Weimer, Y. Yu, R. Ramakrishnan, and C. Curino, “Extending relational query processing with ML inference,” in *10th Conference on Innovative Data Systems Research, CIDR 2020, Amsterdam, The Netherlands, January 12-15, 2020, Online Proceedings*. www.cidrdb.org, 2020. [Online]. Available: <http://cidrdb.org/cidr2020/papers/p24-karanasos-cidr20.pdf>
- [10] J. M. Hellerstein, C. Ré, F. Schoppmann, D. Z. Wang, E. Fratkin, A. Gorajek, K. S. Ng, C. Welton, X. Feng, K. Li, and A. Kumar, “The madlib analytics library: Or mad skills, the sql,” *Proc. VLDB Endow.*, vol. 5, no. 12, p. 1700–1711, aug 2012. [Online]. Available: <https://doi.org/10.14778/2367502.2367510>
- [11] M. Boehm, M. W. Dusenberry, D. Eriksson, A. V. Evfimievski, F. M. Manshadi, N. Parsane, B. Reinwald, F. R. Reiss, P. Sen, A. C. Surve, and S. Tatikonda, “Systemml: Declarative machine learning on spark,” *Proc. VLDB Endow.*, vol. 9, no. 13, p. 1425–1436, sep 2016. [Online]. Available: <https://doi.org/10.14778/3007263.3007279>
- [12] S. Giannakopoulou, M. Karpapothakis, and A. Ailamaki, “Query-driven repair of functional dependency violations,” in *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*. IEEE, 2020, pp. 1886–1889. [Online]. Available: <https://doi.org/10.1109/ICDE48307.2020.00195>
- [13] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” in *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2013. [Online]. Available: <http://arxiv.org/abs/1301.3781>
- [14] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching word vectors with subword information,” *Trans. Assoc. Comput. Linguistics*, vol. 5, pp. 135–146, 2017. [Online]. Available: https://doi.org/10.1162/tacl_a_00051
- [15] “Wikidata.” [Online]. Available: <https://www.wikidata.org/>
- [16] “Common crawl.” [Online]. Available: <https://commoncrawl.org/>
- [17] B. Edizel, A. Piktus, P. Bojanowski, R. Ferreira, E. Grave, and F. Silvestri, “Misspelling oblivious word embeddings,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 3226–3234. [Online]. Available: <https://doi.org/10.18653/v1/n19-1326>
- [18] S. Palkar, J. Thomas, D. Narayanan, P. Thaker, R. Palamuttam, P. Negi, A. Shanbhag, M. Schwarzkopf, H. Pirk, S. P. Amarasinghe, S. Madden, and M. Zaharia, “Evaluating end-to-end optimization for data analytics applications in weld,” *Proc. VLDB Endow.*, vol. 11, no. 9, pp. 1002–1015, 2018. [Online]. Available: <http://www.vldb.org/pvldb/vol11/p1002-palkar.pdf>
- [19] E. Begoli, J. Camacho-Rodríguez, J. Hyde, M. J. Mior, and D. Lemire, “Apache calcite: A foundational framework for optimized query processing over heterogeneous data sources,” in *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, G. Das, C. M. Jermaine, and P. A. Bernstein, Eds. ACM, 2018, pp. 221–230. [Online]. Available: <https://doi.org/10.1145/3183713.3190662>
- [20] J. Johnson, M. Douze, and H. Jégou, “Billion-scale similarity search with GPUs,” *IEEE Transactions on Big Data*, vol. 7, no. 3, pp. 535–547, 2019.
- [21] P. Chrysogelos, “Efficient analytical query processing on cpu-gpu hardware platforms,” *Doctoral Thesis*, p. 132, 2022. [Online]. Available: <http://infoscience.epfl.ch/record/296204>
- [22] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 4171–4186. [Online]. Available: <https://doi.org/10.18653/v1/n19-1423>
- [23] L. J. Orr, S. Kandula, and S. Chaudhuri, “Pushing data-induced predicates through joins in big-data clusters,” *Proc. VLDB Endow.*, vol. 13, no. 3, pp. 252–265, 2019. [Online]. Available: <http://www.vldb.org/pvldb/vol13/p252-orr.pdf>
- [24] S. Hagedorn, S. Kläbe, and K. Sattler, “Putting pandas in a box,” in *11th Conference on Innovative Data Systems Research, CIDR 2021, Virtual Event, January 11-15, 2021, Online Proceedings*. www.cidrdb.org, 2021. [Online]. Available: http://cidrdb.org/cidr2021/papers/cidr2021_paper07.pdf
- [25] N. P. Jouppi, C. Young, N. Patil, D. A. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, R. Boyle, P. Cantin, C. Chao, C. Clark, J. Coriell, M. Daley, M. Dau, J. Dean, B. Gelb, T. V. Ghaemmhami, R. Gottipati, W. Gulland, R. Hagmann, C. R. Ho, D. Hogberg, J. Hu, R. Hundt, D. Hurt, J. Ibarz, A. Jaffey, A. Jaworski, A. Kaplan, H. Khaitan, D. Killebrew, A. Koch, N. Kumar, S. Lacy, J. Laudon, J. Law, D. Le, C. Leary, Z. Liu, K. Lucke, A. Lundin, G. MacKean, A. Maggiore, M. Mahony, K. Miller, R. Nagarajan, R. Narayanaswami, R. Ni, K. Nix, T. Norrie, M. Omernick, N. Penukonda, A. Phelps, J. Ross, M. Ross, A. Salek, E. Samadiani, C. Severn, G. Sizikov, M. Snellman, J. Souter, D. Steinberg, A. Swing, M. Tan, G. Thorson, B. Tian, H. Toma, E. Tuttle, V. Vasudevan, R. Walter, W. Wang, E. Wilcox, and D. H. Yoon, “In-datacenter performance analysis of a tensor processing unit,” in *Proceedings of the 44th Annual International Symposium on Computer Architecture, ISCA 2017, Toronto, ON, Canada, June 24-28, 2017*. ACM, 2017, pp. 1–12. [Online]. Available: <https://doi.org/10.1145/3079856.3080246>

- [26] D. He, S. C. Nakandala, D. Banda, R. Sen, K. Saur, K. Park, C. Curino, J. Camacho-Rodríguez, K. Karanasos, and M. Interlandi, "Query processing on tensor computation runtimes," *Proc. VLDB Endow.*, vol. 15, no. 11, pp. 2811–2825, 2022. [Online]. Available: <https://www.vldb.org/pvldb/vol15/p2811-he.pdf>
- [27] "Wikichip: Neural processor." [Online]. Available: https://en.wikichip.org/wiki/neural_processor
- [28] V. Sanca and A. Ailamaki, "Sampling-based AQP in modern analytical engines," in *International Conference on Management of Data, DaMoN 2022, Philadelphia, PA, USA, 13 June 2022*, S. Blanas and N. May, Eds. ACM, 2022, pp. 4:1–4:8. [Online]. Available: <https://doi.org/10.1145/3533737.3535095>
- [29] P. Sioulas, V. Sanca, I. Mytilinis, and A. Ailamaki, "Accelerating complex analytics using speculation," in *11th Conference on Innovative Data Systems Research, CIDR 2021, Virtual Event, January 11-15, 2021, Online Proceedings*. www.cidrdb.org, 2021. [Online]. Available: http://cidrdb.org/cidr2021/papers/cidr2021_paper03.pdf
- [30] I. Alagiannis, R. Borovica, M. Branco, S. Idreos, and A. Ailamaki, "Nodb: efficient query execution on raw data files," in *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2012, Scottsdale, AZ, USA, May 20-24, 2012*, K. S. Candan, Y. Chen, R. T. Snodgrass, L. Gravano, and A. Fuxman, Eds. ACM, 2012, pp. 241–252. [Online]. Available: <https://doi.org/10.1145/2213836.2213864>
- [31] M. Karpathiotakis, I. Alagiannis, and A. Ailamaki, "Fast queries over heterogeneous data through engine customization," *Proc. VLDB Endow.*, vol. 9, no. 12, pp. 972–983, 2016. [Online]. Available: <http://www.vldb.org/pvldb/vol9/p972-karpathiotakis.pdf>
- [32] J. Wang, X. Yi, R. Guo, H. Jin, P. Xu, S. Li, X. Wang, X. Guo, C. Li, X. Xu, K. Yu, Y. Yuan, Y. Zou, J. Long, Y. Cai, Z. Li, Z. Zhang, Y. Mo, J. Gu, R. Jiang, Y. Wei, and C. Xie, "Milvus: A purpose-built vector data management system," in *SIGMOD '21: International Conference on Management of Data, Virtual Event, China, June 20-25, 2021*, G. Li, Z. Li, S. Idreos, and D. Srivastava, Eds. ACM, 2021, pp. 2614–2627. [Online]. Available: <https://doi.org/10.1145/3448016.3457550>
- [33] J. Choi, Z. Wang, S. Venkataramani, P. I. Chuang, V. Srinivasan, and K. Gopalakrishnan, "PACT: parameterized clipping activation for quantized neural networks," *CoRR*, vol. abs/1805.06085, 2018. [Online]. Available: <http://arxiv.org/abs/1805.06085>
- [34] K. Ramachandra, K. Park, K. V. Emani, A. Halverson, C. A. Galindo-Legaria, and C. Cunningham, "Froid: Optimization of imperative programs in a relational database," *Proc. VLDB Endow.*, vol. 11, no. 4, pp. 432–444, 2017. [Online]. Available: <http://www.vldb.org/pvldb/vol11/p432-ramachandra.pdf>
- [35] A. Herlihy, P. Chrysogelos, and A. Ailamaki, "Boosting efficiency of external pipelines by blurring application boundaries," in *12th Conference on Innovative Data Systems Research, CIDR 2022, Chaminade, CA, USA, January 9-12, 2022*. www.cidrdb.org, 2022. [Online]. Available: <https://www.cidrdb.org/cidr2022/papers/p81-herlihy.pdf>
- [36] M. Abadi, M. Isard, and D. G. Murray, "A computational model for tensorflow: An introduction," in *Proceedings of the 1st ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, ser. MAPL 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 1–7. [Online]. Available: <https://doi.org/10.1145/3088525.3088527>
- [37] A. Y. Halevy, M. J. Franklin, and D. Maier, "Principles of dataspace systems," in *Proceedings of the Twenty-Fifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 26-28, 2006, Chicago, Illinois, USA*, S. Vansummeren, Ed. ACM, 2006, pp. 1–9. [Online]. Available: <https://doi.org/10.1145/1142351.1142352>
- [38] M. J. Franklin, A. Y. Halevy, and D. Maier, "From databases to dataspace: a new abstraction for information management," *SIGMOD Rec.*, vol. 34, no. 4, pp. 27–33, 2005. [Online]. Available: <https://doi.org/10.1145/1107499.1107502>
- [39] I. Burstein, "Nvidia data center processing unit (DPU) architecture," in *IEEE Hot Chips 33 Symposium, HCS 2021, Palo Alto, CA, USA, August 22-24, 2021*. IEEE, 2021, pp. 1–20. [Online]. Available: <https://doi.org/10.1109/HCS52781.2021.9567066>
- [40] T. Neumann, "Efficiently compiling efficient query plans for modern hardware," *Proc. VLDB Endow.*, vol. 4, no. 9, pp. 539–550, 2011. [Online]. Available: <http://www.vldb.org/pvldb/vol4/p539-neumann.pdf>
- [41] A. Pavlo, G. Angulo, J. Arulraj, H. Lin, J. Lin, L. Ma, P. Menon, T. C. Mowry, M. Perron, I. Quah, S. Santurkar, A. Tomasic, S. Toor, D. V. Aken, Z. Wang, Y. Wu, R. Xian, and T. Zhang, "Self-driving database management systems," in *8th Biennial Conference on Innovative Data Systems Research, CIDR 2017, Chaminade, CA, USA, January 8-11, 2017, Online Proceedings*. www.cidrdb.org, 2017. [Online]. Available: <http://cidrdb.org/cidr2017/papers/p42-pavlo-cidr17.pdf>
- [42] P. Chrysogelos, M. Karpathiotakis, R. Appuswamy, and A. Ailamaki, "Hetexchange: Encapsulating heterogeneous CPU-GPU parallelism in JIT compiled engines," *Proc. VLDB Endow.*, vol. 12, no. 5, pp. 544–556, 2019. [Online]. Available: <http://www.vldb.org/pvldb/vol12/p544-chrysogelos.pdf>
- [43] C. Lattner and V. Adve, "Llvm: A compilation framework for lifelong program analysis & transformation," in *International Symposium on Code Generation and Optimization, 2004. CGO 2004*. IEEE, 2004, pp. 75–86.
- [44] S. Idreos, K. Zoumpatianos, B. Hentschel, M. S. Kester, and D. Guo, "The data calculator: Data structure design and cost synthesis from first principles and learned cost models," in *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, G. Das, C. M. Jermaine, and P. A. Bernstein, Eds. ACM, 2018, pp. 535–550. [Online]. Available: <https://doi.org/10.1145/3183713.3199671>
- [45] J. Ding, R. C. Marcus, A. Kipf, V. Nathan, A. Nrusimha, K. Vaidya, A. van Renen, and T. Kraska, "Sagedb: An instance-optimized data analytics system," *Proc. VLDB Endow.*, vol. 15, no. 13, pp. 4062–4078, 2022. [Online]. Available: <https://www.vldb.org/pvldb/vol15/p4062-ding.pdf>
- [46] C. Lattner, M. Amini, U. Bondhugula, A. Cohen, A. Davis, J. Pienaar, R. Riddle, T. Shpeisman, N. Vasilache, and O. Zinenko, "MLIR: Scaling compiler infrastructure for domain specific computation," in *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, 2021, pp. 2–14.