



Quo Vadis ChatGPT? From large language models to Large Knowledge Models

Venkat Venkatasubramanian^{*}, Arijit Chakraborty

Complex Resilient Intelligent Systems Laboratory, Department of Chemical Engineering, Columbia University, New York, NY 10027, USA

ARTICLE INFO

Keywords:

Artificial intelligence
Large language model
Knowledge graph
Domain-specific language processing
Machine learning
Hybrid AI

ABSTRACT

The startling success of ChatGPT and other large language models (LLMs) using transformer-based generative neural network architecture in applications such as natural language processing and image synthesis has many researchers excited about potential opportunities in process systems engineering (PSE). The almost human-like performance of LLMs in these areas is indeed very impressive, surprising, and a major breakthrough. Their capabilities are very useful in certain tasks, such as writing first drafts of documents, code writing assistance, text summarization, etc. However, their success is limited in highly scientific domains as they cannot yet reason, plan, or explain due to their lack of in-depth mechanistic domain knowledge. This is a problem in domains such as chemical engineering as they are governed by fundamental laws of physics and chemistry (and biology), constitutive relations, and highly technical knowledge about materials, processes, and systems. Although purely data-driven machine learning has its immediate uses, the long-term success of AI in scientific and engineering domains would depend on developing hybrid AI systems that combine first principles and technical knowledge effectively. We call these hybrid AI systems *Large Knowledge Models* (LKMs), as they will not be limited to only NLP-based techniques or NLP-like applications. In this paper, we discuss the challenges and opportunities in developing such systems in chemical engineering.

1. Introduction

In recent years, the formidable combination of large data sets, new machine learning (ML) algorithms and architectures, and powerful hardware has led to the birth of large language models (LLMs), such as ChatGPT (Brown et al., 2020), LLaMA (Touvron et al., 2023a,b; AI@Meta, 2024), and Gemini (formerly, Bard) (Team et al., 2023), which have been surprisingly successful in applications such as natural language processing (NLP) and tasks mimicking natural language understanding (NLU). These models are characterized by their vast number of parameters, deep learning capabilities, and extensive training data, which enable them to generate human-like text, comprehend complex instructions, and even perform creative tasks. In simple terms, LLMs are highly sophisticated autocomplete engines that learn to capture conditional probabilistic associations among data elements (called *tokens*) at a massive scale. Although language models are an old idea from the 1980s, it is their recent avatar as LLMs that has caught everyone's attention.

LLMs are based on the transformer architecture introduced by Vaswani et al. in their seminal paper (Vaswani et al., 2017). This architecture enables the model to handle sequential data and understand the context in a more flexible and efficient manner compared to previous models, such as recurrent neural networks (RNNs) or long short-term memory (LSTM) networks. The training process involves

unsupervised learning from a vast corpus of text data, allowing the model to grasp a wide range of language patterns, styles, and information. The massive scale of LLMs is the defining feature. For example, GPT-3 (Generative Pre-trained Transformer-3) developed by Open AI – which is the model behind the chatbot ChatGPT – has over 175 billion parameters and requires approximately 570 GB of filtered text data to train. This immense scale allows the model to develop a broad and nuanced “understanding” of language and context, though it also raises challenges in terms of computational resources and potential biases in the training data. It turns out that what Nobel Laureate Philip Anderson said (Anderson, 1972) more than 50 years ago in the context of physics is valid in AI as well: “More is different”, indeed!

LLMs seem to be good at certain applications, such as writing drafts, coding assistance, summarizing, translating, and answering questions. Despite their capabilities, LLMs face several limitations. They can generate misleading or biased information, relying on the data on which they were trained (Bender et al., 2021). Moreover, ethical considerations about consent, privacy, and misuse of generated content also plague these systems. This was most recently exemplified by the lawsuit filed by The New York Times against OpenAI and Microsoft for copyright infringement in December 2023, in which The New York Times claimed that millions of their articles were used to train AI

^{*} Corresponding author.

E-mail addresses: venkat@columbia.edu (V. Venkatasubramanian), ac4758@columbia.edu (A. Chakraborty).

systems, which are now viewed as a comparable source of information (Grynbaum and Mac, 2023-12-27). In addition to such concerns, the sciences and engineering domains pose other challenges. LLMs are perhaps appropriate for domains such as NLP, where there are no conservation laws or first-principles-based knowledge to leverage. All relevant knowledge is in the data. Therefore, purely probability-based autocomplete techniques are appropriate and successful in these applications.

However, domains such as chemical engineering are governed by the fundamental laws of physics and chemistry (and biology), constitutive relations, and highly technical knowledge about materials, processes, and systems (Venkatasubramanian, 2019b; Venkatasubramanian and Mann, 2022). Not using such a treasure trove of information seems not only inefficient, but also can lead to results that are unsafe. The cost of a mistake in movie and restaurant recommendations made using autocomplete-like guesses is not very high; maybe one loses a couple of hours and a few hundred dollars. However, guessing the wrong decision about the process could lead to potentially dangerous results. Although purely data-driven machine learning has immediate uses, we believe that the long-term success of AI in scientific and engineering domains would depend on the use of first principles and technical knowledge effectively. Current LLMs cannot reason or plan, as they lack such in-depth models of their domains. ChatGPT hallucinations are perhaps interesting, perhaps even amusing, in certain applications, but they can be potentially dangerous in highly technical domains such as chemical engineering.

Furthermore, many chemical engineering applications are not “big data”. We certainly have access to more data now than we did, say, a decade ago. But unlike vision, NLP, and game playing, we cannot easily generate terabytes of data, except perhaps in computer simulations. So, purely data-driven techniques mimicked from such domains are not appropriate. On the other hand, our knowledge of first principles can be leveraged and exploited to reduce the need for large amounts of data. Therefore, developing hybrid AI models is more appropriate for many chemical engineering applications (Venkatasubramanian and Mann, 2022; Mann and Venkatasubramanian, 2021; Mann et al., 2022, 2023a,b, 2024; Chakraborty et al., 2021, 2020). The importance of using domain knowledge has become evident even in non-technical areas for LLMs. For example, ChatGPT uses human experts’ guidance in the last stages of its training as reinforcement learning using human feedback (RLHF). Although this is a step in the right direction, much more needs to be done before the scientific and engineering domains can rely on these tools to solve problems.

In this perspective article, we discuss the history of language models, progress in state-of-the-art systems, and their applications in a few areas. We also highlight the benefits accrued from the use of a hybrid AI approach (Chakraborty et al., 2022) in conjunction with LLMs for a few application domains. Finally, we conclude by suggesting some opportunities for work in the near term in this rapidly evolving field.

2. Evolution of language models

The recent developments in the domain of large language models (LLMs) are groundbreaking in their impact on nearly every aspect of human life. To get a better sense of all this, it is useful to review the early stages of the advances that have enabled such astounding progress. In this section, we present the chronology of language models in three parts by splitting the timeline approximately based on the progress made, acknowledging the capabilities and limitations of that period.

2.1. Early years: The symbolic AI approach – 1950–1990

In the early part of AI’s formative history, the predominant approach for knowledge modeling was based on *symbolic logic*, such as the use of heuristic rules in expert systems, which yielded the name

of that era as *symbolic AI*. This is in contrast with recent developments that are more data-driven, *i.e.*, *numeric AI*, or machine learning (ML) as we know it today.

The fact that language models are not novelties that have emerged only recently during the neural network phase of AI might seem surprising. The conceptualization of chatbot-like systems was made in the 1950s by the computing pioneer Alan Turing, as proposed in his groundbreaking article (Turing, 1950) describing the Turing test. This was quickly followed by the first forays into developing a chatbot, in the format we now know of, named *ELIZA* (Weizenbaum, 1966). Created nearly six decades ago, it is strikingly similar to the chatbots of today in many ways. Its key algorithm relies on scanning the user’s input text for any *keywords*, followed by presenting the output text according to a *rule* related to the keyword(s) identified. If a keyword is not found, then either a context-free statement or a previous output is generated. The progress of early language models in the symbolic AI era is presented in Fig. 1.

It is interesting to note that these rules are not specific to any task or subject and can be modified based on the end-user’s application. *ELIZA* laid the foundation for the development of future chatbots by addressing challenges that must be overcome to advance the field. One such aspect highlighted was the need for such a conversational program to be capable of storing information obtained through its inputs. Subsequently, an information retrieval algorithm would be necessary to improve its conversational capabilities. In addition, such a tool should be able to make inferences based on the input received. This would enable it to interpret and store contextual information described in earlier input(s). These are characteristics that enable the state-of-the-art chatbots today to present themselves as human-like.

Although *ELIZA* presented the first conversational program, it had drawbacks such as its inability to parse the grammar and semantics of its language of use, lack of storage of information received during its input–output interactions, and absence of domain-specific information. In an attempt to address these shortcomings, *SHRDLU* (Winograd, 1971) was created in 1970. It relied on propositional logic, and semantic parsing of grammar rules to provide more flexibility to a conversational program in its learning. This would result in a more human-like interaction with the user. As a proof-of-concept, *SHRDLU* operated as if the user had been conversing with a one-eyed robot with a hand capable of interacting with blocks on a table. This interaction between the user, the robot, and the blocks in this environment demonstrates a deeper understanding of the subject than earlier systems.

Other systems based on the progress of that time were *PARRY* (Colby et al., 1971), a program that responded as if it were a paranoid schizophrenic. The output of the program was created to be a function of its input, beliefs, affects, and intentions (Colby, 1974). It remains one of the first programs capable of dialogue generation. There have been intriguing conversations between *PARRY* (as the patient) and *ELIZA* (as the psychotherapist), which remain one of the few instances of interactions between computers in entirely natural language.

With the increasing success of previous language models, the Hearsay-II Speech Understanding System (Erman et al., 1980) was developed to convert speech to text that a computer program can communicate with. Its innovation came from a hierarchical structure in its organization of linguistic knowledge, which significantly facilitated information retrieval compared to its contemporaries. The system represented an early effort to integrate probabilistic modeling into language processing, utilizing a blend of rule-based and statistical techniques.

It is particularly interesting to note that efforts towards a hybrid approach towards language models and AI in general are not new. Several of the problems that are being addressed today were conceptually formalized in the 1980s and 1990s. *Cyc* (Lenat et al., 1985; Lenat and Marcus, 2023), a natural language system developed by Douglas Lenat nearly four decades ago, is one such tool. Today, it includes over 10 million logic rules coupled with a reasoning engine that is capable

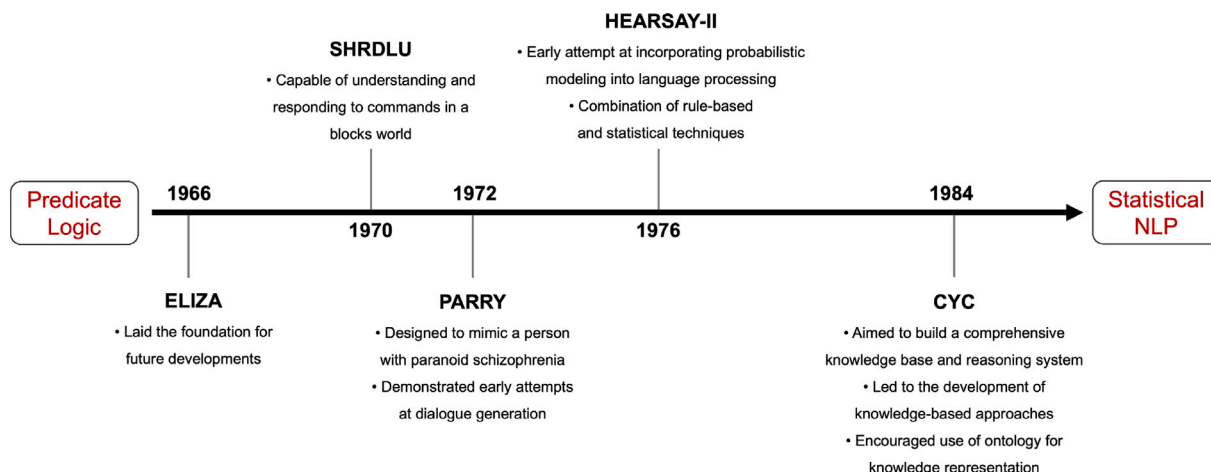


Fig. 1. Progress of early language models, as research transitions from a logic-based approach to statistical-based approach. In chronological order: ELIZA (Weizenbaum, 1966), SHRDLU (Winograd, 1971), PARRY (Colby et al., 1971), HEARSAY (Erman et al., 1980), and CYC (Lenat et al., 1985).

of making inferences. It uses a knowledge graph-like representation with more than 1100 specialized reasoning modules. Dr. Lenat was once quoted as saying, “Intelligence is ten million rules”. (Lenat and Feigenbaum, 1991; Lenat, 1988) This is what his project, Cyc, has now acquired, and it is capable of combining human-like natural language understanding with complex scenarios.

2.2. The data-driven statistical approach – 1990–2017

Despite the efforts of numerous researchers in NLP at the time, the complexities of natural language posed a barrier that required a radically different approach to overcome it, one that can learn from vast amounts of text without the need to be overly prescriptive in its design. With the availability of more data and computational power, largely due to Moore’s law, data-driven statistical techniques made learning from textual sources considerably easier, and more computationally efficient.

One of the first endeavors in statistical language modeling was the use of N-grams (the 1980s), which, combined with Hidden Markov Models (HMMs), led to considerable improvement in a variety of subdomains in NLP such as part-of-speech tagging, speech synthesis, and machine translation. The Viterbi algorithm (Viterbi, 1967) was adopted to decode the most likely sequence of hidden states in HMMs, improving the accuracy of NLP systems. These were followed by statistical methods, which combined information-theoretic approaches. These paved the way for the eventual application of ML-based approaches such as support vector machines (SVMs) in the early 2000s. Such improvements led to significant improvements in accuracy on a multitude of NLP tasks.

The growing number of datasets made available to the research community during this time further catalyzed the growth of statistics-based approaches for NLP tasks. The large-scale annotated corpora, Penn Treebank for English (Marcus et al., 1993), revolutionized the field by providing researchers with standardized datasets to train and evaluate statistical models. These enabled researchers to explore more sophisticated statistical techniques. Soon, the earlier methods of logic-based and rule-based systems would not be able to scale up on these larger datasets.

With the rediscovery of the backpropagation algorithm in 1986 (Rumelhart et al., 1986), the application of neural networks to natural language tasks began to take off. With the continued growth in computational speed and availability of more data than before, several application-specific variants of NNs were proposed. Recurrent neural networks (RNNs) were becoming popular because of their ability to understand the temporal aspect of data, making them ideal for

processing sequential data as found in NLP. Improvements in conventional RNNs were the use of long-short-term-memory (LSTM) networks (Hochreiter and Schmidhuber, 1997; Gers and Schmidhuber, 2001), which emerged as powerful models tailored for sequential data processing. The LSTM models addressed the vanishing gradient problem of RNNs, and were able to account for long-range dependencies, both drawbacks of the conventional RNN architecture. They soon became state-of-the-art models, surpassing their contemporary statistical model counterparts.

2.3. The transformer revolution – 2017-present

The introduction of the transformer architecture (Vaswani et al., 2017) was a watershed moment for the field of NLP, which has since led to stunning progress. A major weakness of previous language models was the inability to retain information from the first elements of a sequence, which was lost when new elements were incorporated into the sequence.

There are 2 blocks in a transformer architecture: the encoder and decoder. The encoder consists of a series of identical layers with each having 2 sub-layers. The first sub-layer performs self-attention, followed by the second sub-layer which is a fully connected feed-forward neural network. The purpose of the encoder block is to derive a contextual representation of the input sequence. The output from the encoder is a vectorized embedding of each item from its input and accounts for its context within the sequence as a result of the self-attention mechanism. This enables the output from the encoder to account for long-range dependencies within a sequence – a significant improvement over prior architectures such as RNNs and LSTM networks. The output of the encoder block serves as the input for the decoder.

The decoder block consists of self-attention, encoder–decoder attention, and feed-forward neural network. The decoder in a transformer model can use the output embedding obtained from the encoder along with the decoder’s outputs to “generate” the next item in the sequence in an autoregressive manner – using the previous output of the decoder as the input for the next item in the sequence – by accounting for relevant context and its own output. Instead of relying on the conventional practice of paying attention only to the last state of the encoder, as is typically done with RNNs, each step of the decoder in the transformer architecture examines all the states of the encoder. This approach enables the decoder to access information pertaining to all elements of the input sequence. This is the strength of state-of-the-art generative pre-trained transformer (GPT) models.

This innovation revolutionized text processing and allowed language models to transition to large language models (LLMs). This can be visualized with the increase in the size of the datasets used for training and the number of parameters in the model, depicted in Fig. 2.

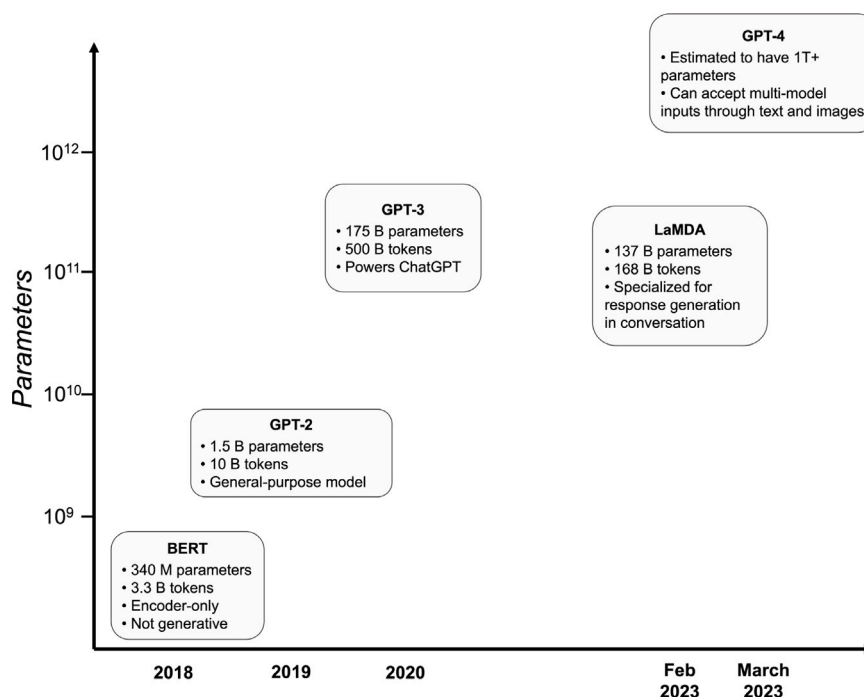


Fig. 2. The rapid increase in the size of LLMs. In chronological order: BERT (Devlin et al., 2018), GPT-2 (Radford et al., 2019), GPT-3 (Brown et al., 2020), LaMDA (Thoppilan et al., 2022), GPT-4 (Achiam et al., 2023). Note the exponential increase in the number of parameters over a few years.

Table 1

The rapid increase in the size of training data used by GPT models by OpenAI. M – million; B – billion; GB – gigabytes.

Year	GPT version	Number of parameters (approximate)	Training corpus size (approximate)
2018	GPT-1	117 M	4.5 GB
2019	GPT-2	1.5 B	40 GB
2020	GPT-3	175 B	570 GB

2.3.1. ChatGPT

In November 2022, OpenAI released ChatGPT, a chatbot that is capable of conversing with a user in a natural language similar to that of a human (Brown et al., 2020). Additionally, it is able to follow instructions based on user prompts, and as a result, it can perform a wide range of tasks – from writing prose and helping solve homework problems to writing code indistinguishable from that written by a professional programmer. It was a landmark moment for LLMs, since its meteoric rise to several million users in a matter of weeks trumped that of any social media platform in human history. This also opened several public discussions on the boons (and banes) of such technologies.

The key improvement in ChatGPT (alternatively, GPT-3/3.5), in contrast to its predecessors, is the fine-tuning approach applied to make ChatGPT more conversational in its usage. This fine-tuning procedure is based on a combination of supervised learning (with labeled data), and reinforcement learning with human feedback (RLHF) (Christiano et al., 2017).

2.3.2. Training data

Since 2018, the number of parameters for each successive GPT model trained has increased exponentially, with a similar increase in the size of the training data (see Table 1). From the conventional ML perspective, this can be viewed as a way to combat the curse of dimensionality. Based on performance, the increase in parameters and training data has yielded more human-like characteristics to the outputs of these models.

Trained on a record-setting 570 GB of training data, GPT-3/3.5 is a successor to previous generative pre-trained transformer (GPT) models, namely GPT-1 (Radford et al., 2018) (2018), followed by GPT-2 (Radford et al., 2019) (2019).

2.3.3. Training procedure

Training for ChatGPT is performed in 3 stages by building on the capabilities of InstructGPT (Ouyang et al., 2022).

Stage 1: Generative pre-training. In the first stage, the model is trained on the training data using conventional NLP techniques such as masked language modeling (MLM), where a portion of the known output is *masked*, and the goal of the trained model is to predict the same. As exemplified by BERT (Devlin et al., 2018), this approach creates a baseline model that is capable of text summarization, translation, and sentiment analysis tasks, among others.

Stage 2: Supervised fine-tuning. Text summarization, translation, and sentiment analysis are one-shot tasks where an input is provided and the model produces an output. Beyond this task, the baseline model is fairly inept at having a conversation with a human on a particular topic, which was the goal of such a chatbot. As a result, it was necessary to fine-tune the baseline model into a supervised fine-tuned (SFT) model, for conversational purposes.

This second stage of training involved humans who created a supervised (labeled) dataset of expected (i.e., considered *ideal*) responses for a set of inputs. These human labelers wrote an appropriate response to the input prompts, to which the baseline model was fine-tuned using stochastic gradient descent (SGD). This is a parameter optimization technique where the dataset is split into multiple smaller batches (mini-batch), and a mini-batch is randomly/stochastically chosen to calculate the gradient of the cost function. The parameters get updated subsequently, and this process is repeated until either a pre-specified value of the cost function is reached, or a pre-specified number of iterations have elapsed, whereby the optimization is terminated. Having obtained a SFT model, one still runs the risk of dealing with an

Table 2

Open-source LLMs are not lagging behind their commercially successful counterparts. K – thousand; M – million; B – billion; T – trillion; GB – gigabyte; TB – terabyte.

Year	GPT version	Number of parameters (approximate)	Training corpus size (approximate)	Number of tokens
2022	BLOOM (Scao et al., 2022)	176 B	1.6 TB	350 B
2023	Llama (Touvron et al., 2023a)	65 B	–	1.4 T
2023	Llama 2 (Touvron et al., 2023b)	70 B	–	2 T
2023	StableLM (A.I., 2023)	3 B & 7 B	825 GB (Gao et al., 2020)	1.5 T
2023	Pythia (Biderman et al., 2023)	70 M - 12 B	825 GB (Gao et al., 2020)	300–334 B
2023	Dolly 2.0 (Conover et al., 2023)	12 B	Fine-tuned on 15 K human-generated prompt/response pairs	Based on Pythia (Biderman et al., 2023)
2023	Alpaca (Taori et al., 2023)	7 B	Fine-tuned on 52K instruction-following data	Fine-tuned from Llama 7B (Touvron et al., 2023a)
2023	Vicuna (Chiang et al., 2023)	13 B	Fine-tuned on 70K user-shared conversations gathered from ShareGPT.com with public APIs	Fine-tuned from Llama (Touvron et al., 2023a)
2024	DBRX (Team, 2024)	132 B	Pretrained on 12 T tokens of carefully curated data, and a maximum context length of 32k tokens	pretrained on 12 T tokens of text and code data
2024	Llama 3 (AI@Meta, 2024)	8 B & 80 B	–	15 T

overtrained model, where the output merely reflects its memorization of the training dataset. This limitation could significantly narrow the scope of application, and diminish the model's efficacy, akin to that of a lookup table.

Stage 3: Reinforcement learning with human feedback (RLHF). In order to address this, a third and final stage of the training involved the use of reinforcement learning (RL), where the model's outputs are rewarded or penalized by *learning* to provide outputs that maximize reward(s). The vanilla implementation of RL (without any modifications) would require a reward function — which is not ideal for chatbots. Further, if such a mathematical function capable of predicting a better output (maximizing reward) existed, it would negate the additional steps being undertaken to fine-tune and improve the model.

Consequently, a human-in-the-loop approach was needed to develop a reward model. Here, human contractors were presented with four to nine outputs from the SFT model, and were asked to rank them from best to the worst responses for the given input prompt. The better the response, the higher the reward; therefore, the reward model learns which kind of response yields a higher reward.

Having discussed the reward model, it is important to update the SFT model so that it overcomes the limitations previously discussed in the earlier stage of training. Here, ChatGPT's model uses Proximal Policy Optimization (PPO) (Schulman et al., 2017), to update the *policy* on the basis of which it returns its outputs. These updates are done in small steps so that the policy does not deviate significantly from what is considered a high-reward output. The information-theoretic concept of Kullback–Leibler (KL) divergence was used to measure such deviations, and to keep the ChatGPT model within reasonable bounds.

2.3.4. Open-source LLMs

The public release of ChatGPT led to a *LLM race* between several AI research laboratories and commercial entities, resulting in the release of several open-source LLMs. These enable the public to tweak the parameters of the pre-trained models, which permit one to apply such a model's complex natural language capabilities to a specific domain.

The most popular open-source LLMs include BLOOM (Scao et al., 2022) and Llama (Touvron et al., 2023a) by the technology giant Meta, which was quickly followed by Llama 2 (Touvron et al., 2023b), and Llama 3 (AI@Meta, 2024). Llama 2, having been trained on 2 trillion tokens, has double the context length of Llama 1. In addition, it offers 3 models of varying sizes (7B, 13B, and 70B parameters) so that the end user can utilize the model most appropriately sized for his or her task. Llama 2 has been trained using a similar approach to that of

ChatGPT by using RLHF. The release of such LLMs has led to a wave of smaller, fine-tuned chatbots that perform comparably to their larger counterparts on specific tasks. These, and numerous other open-source LLMs/chatbots, have been presented in Table 2.

3. Modeling knowledge in chemical engineering: The three paradigms

Understanding the role of AI and LLMs in chemical engineering requires examining it through the lens of various knowledge-modeling approaches employed in our field. Initially, chemical engineering relied heavily on empirical methods and heuristics, with a conspicuous absence of quantitative, first-principles-based models for almost a century (Venkatasubramanian, 2019b). This changed in the 1950s, marked by what is known as the Amundson era (Ramkrishna, 2013), in which the adoption of applied mathematical techniques, especially linear algebra, ordinary differential equations, and partial differential equations, revolutionized the field by enabling the creation of models for unit operations grounded in fundamental principles (Acrivos and Amundson, 1955; Ramkrishna and Amundson, 2004).

Similarly, decision-making within process systems engineering initially relied mainly on empirical knowledge and heuristic approaches. This paradigm shifted in the 1960s, marking another pivotal moment in modeling with the advent of mathematical programming techniques like mixed-integer linear programming (MILP) and mixed-integer nonlinear programming (MINLP) (Daoutidis and Zhang, 2022). This transformation was pioneered by Roger Sargent (Sargent, 1967), and his students.

The next significant avatar in this long evolution of modeling paradigms is the introduction of knowledge representation concepts, and search techniques from artificial intelligence. This started in the early 1980s under the leadership of Westerberg, Stephanopoulos, and others from that era (Banares-Alcantara et al., 1985b, 1987; Stephanopoulos et al., 1987, 1990). After remaining in the background as a fringe activity for the past three decades, pursued by only a few researchers, this knowledge-modeling paradigm has now gone mainstream. The newest avatar in this evolution of knowledge models is the large-language model. While the classical AI models of the 1970s–1980s were based on symbolic representations using logic and reasoning, recent AI advances use statistical machine learning based on probability theory and network science.

Broadly speaking, one might consider the Amundson era as the introduction of formal methods for modeling process units. The Sargent era, and the AI era are about modeling the process engineer.

That is, modeling and automating human information processing and decision-making, formally, to solve problems in synthesis, design, control, scheduling, optimization, and risk analysis. Some of these could be addressed by the mathematical programming framework, *i.e.*, the Sargent approach, but others, such as fault diagnosis and process hazards analysis, require causal model-based reasoning, and are better addressed by AI concepts and techniques.

3.1. Role of symbolic AI in knowledge modeling and representation

When it comes to modeling knowledge, most chemical engineers immediately think of differential and algebraic equations (*i.e.*, DAE models). These are suitable for certain classes of problems, such as those found in thermodynamics, transport phenomena, and reaction engineering.

However, there are other kinds of knowledge that do not lend themselves to such models. For example, reasoning about cause and effect in a process plant is central to fault diagnosis, risk analysis, alarm management, and supervisory control. Knowledge modeling for this problem class typically does not lend itself to the traditional DAE approach to modeling because it cannot provide explicit relationships between cause(s) and effect(s). In some simple cases, perhaps it can, but it is incapable of addressing real-life industrial process systems, which are often complex and nonlinear with incomplete and/or uncertain data. Furthermore, even for simple systems, DAE-based models are not suitable for generating mechanistic explanations of causal behavior. This is where symbolic AI comes in.

We should not forget that the conceptual breakthrough of representing and reasoning with symbolic structures and relationships is an essential contribution of AI (Rich, 1985; Venkatasubramanian, 1989; Ungar and Venkatasubramanian, 1990; Russell and Norvig, 2016). This is what we refer to as symbolic AI, the classical AI of the 1960s-1980s, to differentiate it from data-driven machine learning, which we call numeric AI. Although the importance of symbolic AI has been largely missed in all the current excitement about data-driven machine learning, we expect it to resurface as we go beyond purely data-driven models towards more comprehensive knowledge-based intelligent systems, which are necessary for many applications in chemical engineering.

The symbolic AI methodologies include models such as the following:

- Graph-theoretical models such as signed digraphs used extensively to perform causal reasoning in the identification of abnormal events, diagnosis, and risk analysis (Iri et al., 1979; Vaidhyanathan and Venkatasubramanian, 1995; Maurya et al., 2003,?)
- Petri nets used for modeling discrete event systems (Johnsson and Årzén, 1998; Viswanathan et al., 1998a,b)
- Rule-based production system models used in expert systems for automating higher-order reasoning (Banares-Alcantara et al., 1985a,b, 1987, 1988; Rich and Venkatasubramanian, 1987; Venkatasubramanian and Rich, 1988)
- Semantic network models such as ontologies used in materials discovery and design, domain-specific compilers, etc. (Banares-Alcantara et al. (1985a,b, 1987, 1988), Rich and Venkatasubramanian (1987), Venkatasubramanian and Rich (1988), Aldea et al. (2003), Venkatasubramanian et al. (2006), Marquardt et al. (2010), Hailemariam and Venkatasubramanian (2010a,b)
- Object-oriented models such as agent-based models used in simulating the behavior and decision-making choices of independent, interacting, entities endowed with complex attributes and decision-making powers (Katare and Venkatasubramanian, 2001; Julka et al., 2002)

All these have far-reaching consequences as we begin to develop more comprehensive hybrid-AI systems such as the following in the near future:

- Combining first-principles with data-driven processing (Sundaram et al., 2001; Viswanathan et al., 2002; Ghosh et al., 2003; Chakraborty et al., 2020, 2021)
- Causal models-based explanatory systems (Vaidhyanathan and Venkatasubramanian, 1995; Rich and Venkatasubramanian, 1987; Venkatasubramanian and Rich, 1988; Rich and Venkatasubramanian, 1989; Vedam and Venkatasubramanian, 1999; Zhao et al., 2005a,b)
- Domain-specific knowledge engines (Venkatasubramanian et al., 2006; Caruthers et al., 2003; Katare et al., 2004; Hsu et al., 2008; Suresh et al., 2010a,b)

Thus, we do not view AI methods as simply useful tools to extract patterns from large amounts of data, even though that benefit is very much there, but as a new knowledge modeling paradigm, the next natural evolutionary stage in the long history of developing formal methods: first applied math (*i.e.*, differential and algebraic equations), then operations research (*i.e.*, math programming) and now artificial intelligence. Conceptually, applied math models numerical relationships between variables and parameters, mathematical programming models relationships between constraints, and AI models relationships between symbolic variables and symbolic structures. In the early years, logic was considered the discipline best suited to provide the formal foundations of AI, but recent developments suggest that probability, statistics, and network science are perhaps better suited. The truth might lie in some combination of both, depending on the application. With generative AI poised to make a great impact in chemical engineering (Decardi-Nelson et al., 2024), increasing the accessibility of LLMs, application-specific training datasets, and the inclusion of expert domain knowledge into the model development pipeline, this is an exciting time for research in this field.

4. Applications

The science and engineering community has a myriad of possibilities to take advantage of this new and rapidly advancing technology. In this section, we discuss a few specific examples. It is not meant to be exhaustive, as the field is evolving rapidly, but is only suggestive of the new and exciting possibilities.

4.1. Finetuning of LLMs

The choice of LLMs for users is growing rapidly. Oftentimes, it is preferable to adapt an existing pre-trained LLM to a specific task. This *fine-tuning* of LLMs has become exceedingly popular (and beneficial) in recent years. This fine-tuning is, in essence, transfer learning for LLMs. There are numerous benefits in fine-tuning LLMs: first, one can tailor the output of a general LLM to a specific domain of application; second, it reduces the computational load of training a much larger model with several billion parameters; third, it can reduce the number of trainable parameters for downstream tasks (Hu et al., 2021).

LLMs have been shown to be capable of learning through successive prompts, and thus can learn in context (Wei et al., 2023). Fine-tuning is a more *permanent* exercise in contrast to in-context learning. Here, instead of tuning the trained parameters of the model, one meticulously crafts the inputs/prompts (prompt engineering) provided to the model. The objective is to guide the output to better harmonize with the desired outcome. As more efforts are being made towards customized language models, an increasing variety of prompts can be provided to LLMs. These include prompts based on style of the inputs such as instruction (*e.g.*, *do this action*), completion (*e.g.*, *complete the missing link*), factual (*e.g.*, *provide references for this statement*) etc. These provide background information to the LLM about the expectation of the response. Another variety of prompts that can be provided include keywords, semantic triples, knowledge graphs, relational databases, among others. These show significant promise for domains where there is ample structured and specialized information. The LLMs used in

such domains can leverage this data in conjunction with its inherent natural language capabilities to provide outputs which are relevant and circumvent conventional problems of LLMs such as hallucinations and unavailability of real-time knowledge.

Building on the concept of in-context learning is the idea of *indexing*. This adds an information retrieval functionality to LLMs for extracting data from textual sources. There are numerous benefits accrued from in-context learning, such as increased accuracy resulting in reduced hallucinations and semantic accuracy. The central theme of indexing LLMs is to use the model as a reasoning engine. Depending on the application domain, appropriate knowledge bases can be created, and semantic searches can be made on the same. This approach of combining a search engine with the natural language reasoning capabilities of an LLM was popularized by Lewis et al. (2020) as retrieval-augmented generation (RAG). A major benefit of such an approach is the ability to update the knowledge base with new information. In rapidly evolving fields such as drug discovery, biotechnology, and renewable energy, to name a few, it is vital that their databases be updated in real-time.

4.2. Explanation generation

A major challenge for the sciences and engineering remains the explainability of mathematical models. The types of models obtained after thorough analyses are wide-ranging – from white-box models (derived from first principles) to black-box models (e.g., neural networks), and gray-box models (Von Stosch et al., 2014; Chakraborty et al., 2022) (which rely on combining the principles of both ends of the modeling spectrum). Although there are statistical approaches (Ribeiro et al., 2016; Lundberg and Lee, 2017) that attempt to explain these models, often for an end-user, an *ideal* explanation would be one that is provided in a natural language fashion. In the domain of model discovery (Wilson and Sahinidis, 2017; Chakraborty et al., 2020; Udrescu and Tegmark, 2020; Chakraborty et al., 2021; Jul-Rasmussen et al., 2023, 2024; Chakraborty et al., 2024), it is of prime importance to be able to explain the effects of the features, and their corresponding independent variables on the outputs.

In an attempt to take advantage of the human-like conversational capabilities of ChatGPT, we apply its prowess to generate meaningful sentences that give the illusion of *understanding* in the same way as a human expert would, to the task of explaining mathematical equations and formulae. This is considerably more challenging than asking ChatGPT to write prose, or help with explaining a homework problem. This is because in order to explain the mathematical equation, a relevant knowledge-base is required for the retrieval of information pertinent to the query. Further, despite the abundance of training data provided, ChatGPT does not have information readily available regarding a new equation or mathematical relation that must be described in a natural language manner — especially one that it has never seen during its training process.

Here, we emphasize the need (and success) of providing relevant context when utilizing such LLMs for generating an explanation in natural language. Let us consider a relatively simple mathematical equation that is well known in reaction kinetics, the Arrhenius equation: $k = Ae^{-E_a/RT}$, where k is the specific reaction rate constant, E_a is the activation energy, R is the universal gas constant, and T is the temperature at which the reaction is occurring. For an expert in the field of reaction kinetics, the Arrhenius equation is an empirical relation denoting a negative-inverse exponential dependence on temperature along with an inverse-exponential dependence on the activation energy. Furthermore, such an expert can contrast this empirical relationship with another equation from the same field of study, e.g., the Eyring–Polanyi equation. Knowledge about the background, derivation, success and failure of the equation, details of the variables involved, etc, are essential for an explanation provided by a human.

For a capable program (such as ChatGPT) to produce human-like explanations for mathematical equations and relations, it will require (but is not limited to) the following:

1. Information about the variables involved in the equation
2. Background of the equation (e.g., its derivation)
3. Any causal relationship(s) between the dependent and independent variables
4. Additional dependencies of the parameters involved (e.g., temperature dependence of the specific rate constant as per the Arrhenius equation, in a rate law)

These details are typically found in the scientific literature, such as research articles, specialized textbooks on the subject matter, and industrial manuals, to name a few. The equation is often surrounded by additional metadata, such as references to related sources, equation numbers, derivations, etc. It is thus imperative that when either a human, or an LLM is asked to provide a human-like explanation for such an equation, the relevant context is accounted for since it has pertinent information that can be used effectively for the same.

In conjunction with relevant context, we highlight the importance of prompt engineering for such an application. Pre-trained LLMs do not have all the relevant context for novel equations/mathematical relations that researchers may discover. Additionally, when a query is made to an LLM to describe the model, the relevant context must be provided first. This provides background information, on the basis of which the LLM can attempt to generate an explanation. One such example highlighting the benefit of added context is depicted in Fig. 3. It should be noted that all existing background information does not necessarily need to be provided at the outset. The user may choose to provide information in varying degrees of complexity, either following some existing guidelines provided in relevant textual source(s), or through the human expertise regarding the subject matter.

4.3. LLMs in chemistry

The application of LLMs is not restricted to text-related tasks. In the chemical and biological sciences, many text-based representations exist that can be used to convert a conventional scientific problem into one that can be parsed in a textual format. The most popular choices are the Simplified Molecular Input Line Entry System (SMILES) (Weininger, 1988; Weininger et al., 1989) to represent chemical structures, and FASTA (Lipman and Pearson, 1985; Pearson and Lipman, 1988) commonly used in bioinformatics. LLMs, with their textual reasoning capabilities, can be taken advantage of for such tasks. Guo et al. (2023) used a variety of GPT models for eight chemistry tasks. These include forward reaction and retrosynthetic prediction, among others, which have been predicted using domain-specific ML approaches. The results show its ability to perform well on text-related explanation tasks such as molecular captioning and text-based molecular design. Unsurprisingly, the results are not much better than the current state-of-the-art non-LLM approaches for the same tasks. This can be linked to an inherent lack of knowledge in the chemistry domain that would be essential – say, for a human expert – given the same tasks.

Consequently, an effort was made to enhance the capabilities of LLMs for chemistry-related tasks, titled ChemCrow (Bran et al., 2023), by integrating helper “tools” that work with the existing natural language reasoning capabilities and enable a wider range of actions. These include, but are not limited to, a web search tool, a literature search tool, a Python coding tool, a molecule manipulation tool (which can alter SMILES, and the corresponding molecules for subsequent tasks), etc. These improvements substantially improve the performance of the LLM and can be used to systematically obtain more relevant results for chemistry-related queries. Furthermore, this can allow the LLM to provide a reason for its steps taken towards obtaining the results, which can greatly benefit a human expert and can provide an additional avenue to explore. Although ChemCrow can have numerous advantages for the scientific community, it is necessary to have safety measures in place to reduce the risk of misuse of the same. Consequently, such tools were added to the framework. Further, for the safety of the user,

You
Explain the following equation:

$$t_t = t_f + t_e + t_c + t_R$$

ChatGPT
This equation seems to represent the total time (t_t) it takes to complete a task, which is comprised of different components:

- t_f : Time for forming a plan or strategy.
- t_e : Time for executing the task.
- t_c : Time spent on communication or coordination related to the task.
- t_R : Time for reviewing or reflecting on the task after completion.

So, the total time to complete the task includes the time spent on planning, execution, communication, and reflection. This equation helps in understanding and analyzing the different factors contributing to the overall time required for a task.

You
Explain the following equation in chemical engineering:

$$t_t = t_f + t_e + t_c + t_R$$

given that
 t_t is total cycle time,
 t_f is time necessary to fill the reactor,
 t_e is time necessary to heat the reactor,
 t_c is time necessary to clean the reactor, and
 t_R is total reaction time.

Context

ChatGPT
In chemical engineering, the equation you provided represents the total cycle time t_t required for a particular reactor operation. Let's break down each term:

1. t_f is the time necessary to fill the reactor. This includes the time it takes to introduce the reactants into the reactor until it's completely filled.
2. t_e is the time necessary to heat the reactor. Some reactions require specific temperatures for optimal performance or to initiate the reaction itself. This term encompasses the time taken to reach and stabilize at the desired temperature.
3. t_c is the time necessary to clean the reactor. After a reaction cycle completes, it's often essential to clean the reactor thoroughly before the next cycle to prevent contamination or undesirable reactions.
4. t_R is the total reaction time. This refers to the duration required for the actual chemical reaction to occur inside the reactor. It includes the time from the start of the reaction until it reaches completion or equilibrium.

The total cycle time t_t is the sum of these individual components t_f , t_e , t_c , and t_R . It represents the overall time required to complete one full cycle of operation in the reactor, from filling to cleaning, including the reaction time itself. This equation helps in the analysis and optimization of reactor processes by considering the various time components involved in the cycle.

Fig. 3. An example of how background context helps guide ChatGPT towards a better explanation of a simple equation from the domain of chemical reaction engineering. The additional context provided, is shown in the red colored box. This exemplifies the use of prompt engineering, which enables a more sensible output from LLMs, especially for more focused question-and-answering tasks such as the one highlighted here.

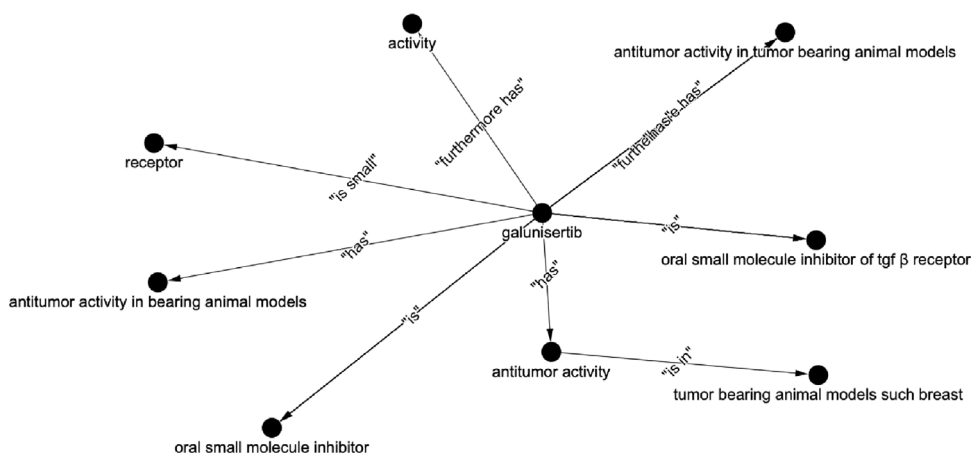


Fig. 4. The output from SUSIE (Mann et al., 2023b), for an abstract in pharmaceutical sciences, discussing galunisertib (Herbertz et al., 2015).

the data is queried from the PubChem database and are provided in a simple manner to the user.

Similarly, graph neural nets that have traditionally been unable to use textual information and lack sufficient training data can be used in conjunction with language models for reaction prediction, as demonstrated by ReLM (Shi et al., 2023). Here, reaction prediction is performed using a pre-trained language model and graph neural networks. Alternate approaches include using the capabilities of natural language in the design of organic structures, as was demonstrated by Ito et al. (2024). Here, GPT-4 (Achiam et al., 2023) was used as a natural language medium to obtain additional insights in an iterative manner.

4.4. LLMs in biology, pharmaceuticals, and drug discovery

Large language models have also become useful in the pharmaceutical domain. A recent example of this is the extraction of semantic triples

performed on pharmaceutical documents by Schema-based Unsupervised Semantic Information Extraction (SUSIE) (Mann et al., 2023b). Such a tool can accelerate drug discovery by searching for keywords and phrases through a corpus of documents in a fraction of the time that it would take a human expert. This system relied on fine-tuning a pre-trained domain-specific language model (BioBERT (Lee et al., 2020)), on data from its intended application domain, followed by information extraction. Such a tool can drastically reduce the time taken by a human to undertake repetitive and tedious manual tasks, to mere seconds. A sample output from SUSIE based on an abstract of a journal article in the pharmaceutical sciences domain is depicted in Fig. 4. When this output – which is obtained after careful inclusion of domain-knowledge expertise into the data-driven pipeline – is provided to an LLM (here, ChatGPT), it is capable of rendering an accurate natural language summary of the text. The same is shown in Fig. 5. Relying on prompt engineering and data augmentation from updated

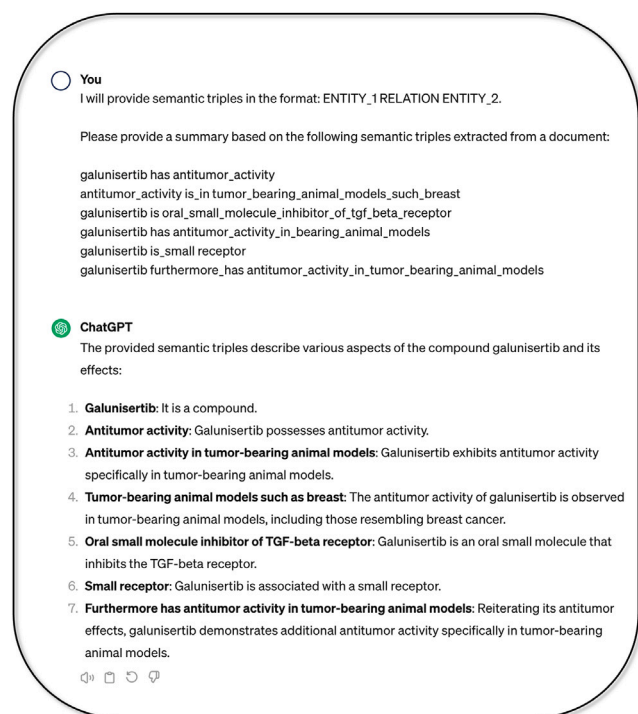


Fig. 5. The output from ChatGPT for a summarization task when provided with a knowledge graph generated by SUSIE (Mann et al., 2023b), for an abstract in pharmaceutical sciences, discussing galunisertib (Herbertz et al., 2015). The summarization of complex scientific text has become simpler, and more accurate by virtue of semantic triples extracted.

sources, these extracted and domain-aware knowledge graphs serve as guardrails for the LLMs and guide the model to yield sensible outputs (Taskiran et al., 2024b,a; Kadiresan et al., 2024).

Other applications in this domain include ChatPathway (Li et al., 2023), which highlights the successful merging of domain knowledge in biology with the power of LLMs to predict biochemical reactions and their pathways. With the rapid increase in the number of applications of LLMs in the scientific and engineering domains, there has been a rise in scientific LLMs such as PubMedBERT (Gu et al., 2021), MolGPT (Bagal et al., 2021), SciBERT (Beltagy et al., 2019) – to name a few. For a more in-depth review of scientific LLMs, the reader is advised to refer to this comprehensive article by Zhang et al. (2024) discussing the same in great detail.

5. Next steps: From LLMs to LKMs

With the advent of the LLMs, Pandora’s box has been opened, revealing many exciting possibilities and serious concerns. As Venkatasubramanian argues (Venkatasubramanian, 2024), ultra-large language models, such as GPT-3.5 and beyond, are not mere autocomplete engines or stochastic parrots. They have new emergent capabilities that require creating a new conceptual framework similar to the transition from Newtonian to statistical mechanics. However, despite their impressive capabilities in specific applications, current LLMs have limited success in highly scientific and engineering applications. This is mainly due to the lack of in-depth domain knowledge based on first principles. Current LLMs, in general, appear to be a mile wide in their scope but only a foot deep in fundamental principles.

This limitation of purely data science models is due to their lack of “understanding” of the underlying knowledge. For example, a self-driving car can navigate impressively through traffic, but does it “know” and “understand” the concepts of mass, momentum, acceleration, force, and Newton’s laws as we do? It does not. Its behavior

is similar to that of a cheetah chasing an antelope in the wild. Both animals show great mastery of the dynamics of the chase, but do they “understand” these concepts? At best, current AI systems have perhaps achieved animal-like mastery of their tasks but have not gained a deeper “understanding” as many humans do. This is a serious problem.

What we need are domain-specific LLMs that are perhaps a foot wide but a mile deep in domain knowledge. For example, a custom LLM in pharmaceutical engineering might not compose poems like Shakespeare, but it will not hallucinate on pharmaceutical chemistry. We need many such custom LLMs in different application domains that can communicate with one another.

This raises the importance of building first-principles knowledge in LLMs. For four decades, the first author has argued (Rich and Venkatasubramanian, 1987, 1989) about the importance of such hybrid AI models, combining mechanistic understanding based on the underlying physics, chemistry, and/or biology of our systems and processes with data-driven techniques. Our domain, in fact, most science and engineering, is governed by such fundamental principles. In this regard, it is different from areas such as game playing and computer vision, where there are no such conservation laws or constitutive equations to exploit. So, it makes sense in those domains to be primarily data-driven. But our applications are quite different.

A related point is that many chemical engineering applications are not “big data”. We certainly have access to more data now than we did, say, a decade ago. But unlike vision or game-playing, we do not generate terabytes of data easily, except perhaps in computer simulations. So, purely data-driven techniques mimicked from such domains are not appropriate and might not work. On the other hand, our first-principles knowledge can be leveraged and exploited to reduce the need for large amounts of data. Thus, the construction of hybrid domain-specific LLMs is more appropriate for many chemical engineering applications.

We classify AI opportunities and challenges into four categories (Venkatasubramanian, 2019a; AIChE, 2019) – “easy”, “hard”, “harder”, and “hardest”, problems. The relatively “easy” ones are those where plenty of data is available, and many standard machine learning software could be used to extract useful patterns from such data. These are not particularly difficult problems, as these were demonstrated even in the 1990s. On the other hand, “hard” problems require the construction of hybrid AI models and the generation of mechanism-based causal explanations. Although work in the 1990s, again, showed (Psychogios and Ungar, 1992) how these could be done, we are still quite a ways from developing such models systematically, reliably, and easily for large-scale and diverse applications. We think this might take another five years or so.

The third category, the “harder” problems, is those in which one would have to build domain-specific ontologies, compilers, languages, etc., such as those reported for catalyst design (Caruthers et al., 2003; Katere et al., 2004) and pharmaceutical manufacturing (Venkatasubramanian et al., 2006; Hailemariam and Venkatasubramanian, 2010a,b). Custom domain-specific LLMs fall into this category (Mann et al., 2023b). This would require much more work than simply applying standard techniques and software in machine learning. They need careful integration of symbolic domain knowledge with data-driven methods. In fact, even generic LLMs such as ChatGPT have realized the importance of human experts and are already doing this under the guise of reinforcement learning with human feedback (RLHF). Another example of this realization is the development of *AlphaGeometry*, which proves mathematical theorems at the Olympiad level (Trinh et al., 2024). As a neuro-symbolic system, it uses a language model in conjunction with a symbolic deduction engine. In another recent example, researchers at Apple (Mirzadeh et al., 2024) observe the illusory and brittle nature of reasoning in LLMs. The domain-specific LLMs that we envision would require much more human expert guidance in the form of ontologies and heuristics, reflecting some of the techniques used in the expert systems

era (Venkatasubramanian, 2019b). Automation tools will certainly come along to make building such ontologies somewhat easier.

There are encouraging results on this front in recent work (Roose, 2024-05-21). With the use of dictionary learning, researchers at Anthropic AI were able to investigate the underlying features (Templeton et al., 2024) of their LLM named Claude Sonnet (Roose, 2023-02-16). This is certainly a step in the right direction, as it enables researchers to better understand the shortcomings of these complex (to date, uninterpretable) models and aids in the future development of more explainable LLMs. Additional efforts at the frontiers of LLM research are geared towards developing small language models (SLMs), which are a fraction of the size of LLMs, are trained on higher quality data, and require significantly lower training computations than LLMs. One such example is that of Microsoft's Phi-1 SLM (Gunasekar et al., 2023), which was trained on a textbook-quality dataset with a size of 7 billion tokens. Building on the success of the earlier model, Microsoft launched the Phi-3-mini SLM (Abdin et al., 2024). These compact yet effective language models are more optimized versions of their larger predecessors.

Along these lines, an exciting, while nascent, avenue to explore is the inclusion of knowledge graphs (KGs) in an LLM framework. This enables one to circumvent the problems of hallucinations in LLMs (Ji et al., 2023), and augment domain-specific knowledge to LLMs, resulting in a system that does not merely recall (occasionally, incorrect) factual information from its training data (Petroni et al., 2019). Due to the inherent lack of interpretability of an LLM, which functions as a black-box model due to its reasoning aided by a probability model (Zhang et al., 2022), it is necessary to incorporate structured knowledge into the same, targeted for a specific domain. For a more comprehensive treatment of the use of structured knowledge representation in KGs and enhanced inference and interpretability for LLMs, the reader is advised to refer to the review article by Pan et al. (2023).

We call these hybrid AI systems *Large Knowledge Models* (LKMs) because they will not be limited to NLP-based techniques or NLP-like applications only. We estimate that such systems might take about 5–10 years to emerge as routine implementations.

Finally, the “hardest” challenge is, as Venkatasubramanian describes (Venkatasubramanian, 2024) in detail, in developing a “Thermodynamics”-like theory of the LLMs. In our opinion, the most interesting and intellectually challenging problems lie in the “hard”, “harder”, and “hardest” categories. We think these are the areas where the PSE research community should focus now.

In conclusion, generative AI has opened up unimagined possibilities in all aspects of human endeavor. It is poised to have a great impact on process systems engineering. However, to harness its potential safely and effectively, we need to go beyond large language models (LLMs) to large knowledge models (LKMs), which incorporate fundamental knowledge and human expertise deeply and effectively.

CRedit authorship contribution statement

Venkat Venkatasubramanian: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Arijit Chakraborty:** Writing – review & editing, Writing – original draft.

Declaration of competing interest

We, the authors, declare that there is no conflict of interest with this submission.

Acknowledgments

This paper is based on the plenary lecture delivered by the first author at the European Symposium of Computer-Aided Process Engineering (ESCAPE 33) in Athens, Greece, in June 2023. VV is grateful for the kind invitation and for the warm hospitality of Professor Antonis Kokossis and his colleagues at the conference. This work is supported in part by the NSF EFRI-DChEM, USA 2132142 grant and funding from the Center for the Management of Systemic Risk (CMSR) at Columbia University, USA.

Data availability

No data was used for the research described in the article.

References

- Abdin, M., Jacobs, S.A., Awan, A.A., Aneja, J., Awadallah, A., Awadalla, H., Bach, N., Bahree, A., Bakhtiari, A., Behl, H., et al., 2024. Phi-3 technical report: A highly capable language model locally on your phone. arXiv preprint arXiv:2404.14219.
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al., 2023. Gpt-4 technical report. arXiv preprint arXiv:2303.08774.
- Acrivios, A., Amundson, N.R., 1955. Applications of matrix mathematics to chemical engineering problems. *Ind. Eng. Chem.* 47 (8), 1533–1541.
- A.I., S., 2023. Stability AI launches the first of its stablelm suite of language models.
- AICHE, 2019. Venkat venkatasubramanian on artificial intelligence in chemical engineering.
- AI@Meta, 2024. Llama 3 model card.
- Aldea, A., Bañares-Alcántara, R., Bocio, J., Gramajo, J., Isern, D., Kokossis, A.C., Jiménez, L., Moreno, A., Riaño, D., 2003. An ontology-based knowledge management platform. In: *IIWeb*. pp. 177–182.
- Anderson, P.W., 1972. More is different. *Science* 177 (4047), 393–396.
- Bagal, V., Aggarwal, R., Vinod, P., Priyakumar, U.D., 2021. MolGPT: molecular generation using a transformer-decoder model. *J. Chem. Inf. Model.* 62 (9), 2064–2076.
- Banares-Alcantara, R., Ko, E., Westerberg, A., Rychener, M., 1988. DECADE—a hybrid expert system for catalyst selection—II. Final architecture and results. *Comput. Chem. Eng.* 12 (9–10), 923–938.
- Banares-Alcantara, R., Sriram, D., Venkatasubramanian, V., Westerberg, A., Rychener, M., 1985a. Knowledge-based expert systems for CAD. *Chem. Eng. Prog.* 81 (9), 25–30.
- Banares-Alcantara, R., Westerberg, A., Ko, E., Rychener, M., 1987. Decade—A hybrid expert system for catalyst selection—I. Expert system consideration. *Comput. Chem. Eng.* 11 (3), 265–277.
- Banares-Alcantara, R., Westerberg, A.W., Rychener, M.D., 1985b. Development of an expert system for physical property predictions. *Comput. Chem. Eng.* 9 (2), 127–142.
- Beltagy, I., Lo, K., Cohan, A., 2019. SciBERT: A pretrained language model for scientific text. arXiv preprint arXiv:1903.10676.
- Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S., 2021. On the dangers of stochastic parrots: Can language models be too big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. pp. 610–623.
- Biderman, S., Schoelkopf, H., Anthony, Q.G., Bradley, H., O'Brien, K., Hallahan, E., Khan, M.A., Purohit, S., Prashanth, U.S., Raff, E., et al., 2023. Pythia: A suite for analyzing large language models across training and scaling. In: *International Conference on Machine Learning*. PMLR, pp. 2397–2430.
- Bran, A.M., Cox, S., White, A.D., Schwaller, P., 2023. ChemCrow: Augmenting large-language models with chemistry tools. arXiv preprint arXiv:2304.05376.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al., 2020. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* 33, 1877–1901.
- Caruthers, J., Lauterbach, J.A., Thomson, K., Venkatasubramanian, V., Snively, C., Bhan, A., Katare, S., Oskarsdottir, G., 2003. Catalyst design: knowledge extraction from high-throughput experimentation. *J. Catalysis* 216 (1–2), 98–109.
- Chakraborty, A., Gandhi, A., Hasan, M.F., Venkatasubramanian, V., 2024. Discovering zeolite adsorption isotherms: a hybrid AI modeling approach. In: *Computer Aided Chemical Engineering*. vol. 53, Elsevier, pp. 511–516.
- Chakraborty, A., Serneels, S., Claussen, H., Venkatasubramanian, V., 2022. Hybrid AI models in chemical engineering—A purpose-driven perspective. *Comput. Aided Chem. Eng.* 51, 1507–1512.
- Chakraborty, A., Sivaram, A., Samavedham, L., Venkatasubramanian, V., 2020. Mechanism discovery and model identification using genetic feature extraction and statistical testing. *Comput. Chem. Eng.* 140, 106900.

- Chakraborty, A., Sivaram, A., Venkatasubramanian, V., 2021. AI-DARWIN: A first principles-based model discovery engine using machine learning. *Comput. Chem. Eng.* 154, 107470.
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J.E., Stoica, I., Xing, E.P., 2023. Vicuna: An open-source chatbot impressing GPT-4 with 90%* ChatGPT quality.
- Christiano, P.F., Leike, J., Brown, T., Martic, M., Legg, S., Amodei, D., 2017. Deep reinforcement learning from human preferences. *Adv. Neural Inf. Process. Syst.* 30.
- Colby, K.M., 1974. Ten criticisms of parry. *ACM SIGART Bull.* (48), 5–9.
- Colby, K.M., Weber, S., Hilf, F.D., 1971. Artificial paranoia. *Art. Intell.* 2 (1), 1–25.
- Conover, M., Hayes, M., Mathur, A., Xie, J., Wan, J., Shah, S., Ghodsi, A., Wendell, P., Zaharia, M., Xin, R., 2023. Free dolly: Introducing the world's first truly open instruction-tuned LLM.
- Daoutidis, P., Zhang, Q., 2022. From Amundson, Aris, and Sargent to the future of process systems engineering. *Chem. Eng. Res. Des.* 188, 704–713.
- Decardi-Nelson, B., Alshehri, A.S., Ajaqekar, A., You, F., 2024. Generative AI and process systems engineering: The next frontier. *arXiv preprint arXiv:2402.10977*.
- Devlin, J., Chang, M.-W., Lee, K., Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Erman, L.D., Hayes-Roth, F., Lesser, V.R., Reddy, D.R., 1980. The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty. *ACM Comput. Surv.* 12 (2), 213–253.
- Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., Presser, S., Leahy, C., 2020. The pile: An 800 GB dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Gers, F.A., Schmidhuber, E., 2001. LSTM recurrent networks learn simple context-free and context-sensitive languages. *IEEE Trans. Neural Netw.* 12 (6), 1333–1340.
- Ghosh, P., Katare, S., Patkar, P., Caruthers, J.M., Venkatasubramanian, V., Walker, K.A., 2003. Sulfur vulcanization of natural rubber for benzothiazole accelerated formulations: from reaction mechanisms to a rational kinetic model. *Rubber Chem. Technol.* 76 (3), 592–693.
- Grynbaum, M.M., Mac, R., 2023-12-27. The times sues openai and microsoft over A.I. Use of copyrighted work. *N.Y. Times*.
- Gu, Y., Tinn, R., Cheng, H., Lucas, M., Usuyama, N., Liu, X., Naumann, T., Gao, J., Poon, H., 2021. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Healthc. (HEALTH)* 3 (1), 1–23.
- Gunasekar, S., Zhang, Y., Aneja, J., Cesar, C., Mendes, T., Giorno, A.D., Gopi, S., Javaheripi, M., Kauffmann, P., de Rosa, G., Saarikivi, O., Salim, A., Shah, S., Singh Behl, H., Wang, X., Bubeck, S., Eldan, R., Kalai, A.T., Lee, Y.T., Li, Y., 2023. Textbooks are all you need.
- Guo, T., Guo, K., Liang, Z., Guo, Z., Chawla, N.V., Wiest, O., Zhang, X., et al., 2023. What indeed can GPT models do in chemistry? A comprehensive benchmark on eight tasks. *arXiv preprint arXiv:2305.18365*.
- Hailamariam, L., Venkatasubramanian, V., 2010a. Purdue ontology for pharmaceutical engineering: part I. Conceptual framework. *J. Pharm. Innov.* 5, 88–99.
- Hailamariam, L., Venkatasubramanian, V., 2010b. Purdue ontology for pharmaceutical engineering: Part II. Applications. *J. Pharm. Innov.* 5, 139–146.
- Herbertz, S., Sawyer, J.S., Stauber, A.J., Gueorguieva, I., Driscoll, K.E., Estrem, S.T., Cleverly, A.L., Desai, D., Guba, S.C., Benhadji, K.A., et al., 2015. Clinical development of galunisertib (LY2157299 monohydrate), a small molecule inhibitor of transforming growth factor-beta signaling pathway. *Drug Design, Development and Therapy* 4479–4499.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.
- Hsu, S.-H., Krishnamurthy, B., Rao, P., Zhao, C., Jagannathan, S., Venkatasubramanian, V., 2008. A domain-specific compiler theory based framework for automated reaction network generation. *Comput. Chem. Eng.* 32 (10), 2455–2470.
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Iri, M., Aoki, K., O'Shima, E., Matsuyama, H., 1979. An algorithm for diagnosis of system failures in the chemical process. *Comput. Chem. Eng.* 3 (1–4), 489–493.
- Ito, S., Muraoka, K., Nakayama, A., 2024. De novo design of organic structure-directing agents for zeolites using a general-purpose large language model. *ChemRxiv*.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P., 2023. Survey of hallucination in natural language generation. *ACM Comput. Surv.* 55 (12), 1–38.
- Johnsson, C., Årzén, K.-E., 1998. Grafchart and its relations to grafset and Petri nets. *IFAC Proc. Vol.* 31 (15), 95–100.
- Jul-Rasmussen, P., Chakraborty, A., Venkatasubramanian, V., Liang, X., Huusom, J.K., 2023. Identifying first-principles models for bubble column aeration using machine learning. In: *Computer Aided Chemical Engineering*. vol. 52, Elsevier, pp. 1089–1094.
- Jul-Rasmussen, P., Chakraborty, A., Venkatasubramanian, V., Liang, X., Huusom, J.K., 2024. Hybrid AI modeling techniques for pilot scale bubble column aeration: A comparative study. *Comput. Chem. Eng.* 108655.
- Julka, N., Srinivasan, R., Karimi, I., 2002. Agent-based supply chain management—I: framework. *Comput. Chem. Eng.* 26 (12), 1755–1769.
- Kadiresan, A., Taskiran, N.P., Chakraborty, A., Mann, V., Venkatasubramanian, V., 2024. Contrastive learning to improve pharmaceutical knowledge graph quality in machine learning. In: *2024 AIChE Annual Meeting*. AIChE.
- Katare, S., Caruthers, J.M., Delgass, W.N., Venkatasubramanian, V., 2004. An intelligent system for reaction kinetic modeling and catalyst design. *Ind. Eng. Chem. Res.* 43 (14), 3484–3512.
- Katare, S., Venkatasubramanian, V., 2001. An agent-based learning framework for modeling microbial growth. *Eng. Appl. Artif. Intell.* 14 (6), 715–726.
- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H., Kang, J., 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 36 (4), 1234–1240.
- Lenat, D.B., 1988. The case for inelligence. In: *Proceedings of the International Workshop on Artificial Intelligence for Industrial Applications*, Tokyo.
- Lenat, D.B., Feigenbaum, E.A., 1991. On the thresholds of knowledge. *Artificial Intelligence* 47 (1), 185–250.
- Lenat, D., Marcus, G., 2023. Getting from generative ai to trustworthy ai: What llms might learn from cyc. *arXiv preprint arXiv:2308.04445*.
- Lenat, D.B., Prakash, M., Shepherd, M., 1985. CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Mag.* 6 (4), 65.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al., 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Adv. Neural Inf. Process. Syst.* 33, 9459–9474.
- Li, Y., Xu, H., Zhao, H., Guo, H., Liu, S., 2023. Chatpathway: Conversational large language models for biology pathway detection. In: *NeurIPS 2023 AI for Science Workshop*.
- Lipman, D.J., Pearson, W.R., 1985. Rapid and sensitive protein similarity searches. *Science* 227 (4693), 1435–1441.
- Lundberg, S.M., Lee, S.-I., 2017. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* 30.
- Mann, V., Brito, K., Gani, R., Venkatasubramanian, V., 2022. Hybrid, interpretable machine learning for thermodynamic property estimation using grammar2vec for molecular representation. *Fluid Phase Equilib.* 561, 113531.
- Mann, V., Gani, R., Venkatasubramanian, V., 2023a. Group contribution-based property modeling for chemical product design: A perspective in the AI era. *Fluid Phase Equilib.* 568, 113734.
- Mann, V., Sales-Cruz, M., Gani, R., Venkatasubramanian, V., 2024. eSFILES: Intelligent process flowsheet synthesis using process knowledge, symbolic AI, and machine learning. *Comput. Chem. Eng.* 181, 108505.
- Mann, V., Venkatasubramanian, V., 2021. Predicting chemical reaction outcomes: A grammar ontology-based transformer framework. *AIChE J.* 67 (3), e17190.
- Mann, V., Viswanath, S., Vaidyaraman, S., Balakrishnan, J., Venkatasubramanian, V., 2023b. SUSIE: Pharmaceutical CMC ontology-based information extraction for drug development using machine learning. *Comput. Chem. Eng.* 179, 108446.
- Marcus, M., Santorini, B., Marcinkiewicz, M.A., 1993. Building a large annotated corpus of English: The Penn Treebank. *Comput. Linguist.* 19 (2), 313–330.
- Marquardt, W., Morbach, J., Wiesner, A., Yang, A., CAPE, O., 2010. *A Re-Usable Ontology for Chemical Process Engineering*. Springer.
- Maurya, M.R., Rengaswamy, R., Venkatasubramanian, V., 2003. A systematic framework for the development and analysis of signed digraphs for chemical processes. 1. Algorithms and analysis. *Ind. Eng. Chem. Res.* 42 (20), 4789–4810.
- Mirzadeh, I., Alizadeh, K., Shahrokhi, H., Tuzel, O., Bengio, S., Farajtabar, M., 2024. GSM-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv preprint arXiv:2410.05229*.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al., 2022. Training language models to follow instructions with human feedback. *Adv. Neural Inf. Process. Syst.* 35, 27730–27744.
- Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., Wu, X., 2023. Unifying large language models and knowledge graphs: A roadmap. *arXiv preprint arXiv:2306.08302*.
- Pearson, W.R., Lipman, D.J., 1988. Improved tools for biological sequence comparison. *Proc. Natl. Acad. Sci.* 85 (8), 2444–2448.
- Petroni, F., Rocktäschel, T., Lewis, P., Bakhtin, A., Wu, Y., Miller, A.H., Riedel, S., 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- Psichogios, D.C., Ungar, L.H., 1992. A hybrid neural network-first principles approach to process modeling. *AIChE J.* 38 (10), 1499–1511.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al., 2018. Improving language understanding by generative pre-training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al., 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1 (8), 9.
- Ramkrishna, D., 2013. The Neal Amundson era. *Rapid evolution of chemical engineering science*. *AIChE J.* 59 (9), 3147–3157.
- Ramkrishna, D., Amundson, N.R., 2004. Mathematics in chemical engineering: A 50 year introspection. *AIChE J.* 50 (1), 7–23.
- Ribeiro, M.T., 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*, pp. 1135–1144.
- Rich, E., 1985. Artificial intelligence and the humanities. *Comput. Humanit.* 19 (2), 117–122.
- Rich, S.H., Venkatasubramanian, V., 1987. Model-based reasoning in diagnostic expert systems for chemical process plants. *Comput. Chem. Eng.* 11 (2), 111–122.

- Rich, S.H., Venkatasubramanian, V., 1989. Causality-based failure-driven learning in diagnostic expert systems. *AIChE J.* 35 (6), 943–950.
- Roose, K., 2023-02-16. A conversation with Bing's chatbot left me deeply unsettled. *N.Y. Times*.
- Roose, K., 2024-05-21. A.I.'s black boxes just got a little less mysterious. *N.Y. Times*.
- Rumelhart, D.E., Hinton, G.E., Williams, R.J., 1986. Learning representations by back-propagating errors. *Nature* 323 (6088), 533–536.
- Russell, S.J., Norvig, P., 2016. *Artificial Intelligence: A Modern Approach*. Pearson.
- Sargent, R., 1967. Integrated design and optimization of processes. *Chem. Eng. Prog.* 63 (9), 71–74.
- Scao, T.L., Fan, A., Akiki, C., Pavlick, E., Ilić, S., Hesslow, D., Castagné, R., Luccioni, A.S., Yvon, F., Gallé, M., et al., 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shi, Y., Zhang, A., Zhang, E., Liu, Z., Wang, X., 2023. ReIm: Leveraging language models for enhanced chemical reaction prediction. *arXiv preprint arXiv:2310.13590*.
- Stephanopoulos, G., Henning, G., Leone, H., 1990. MODEL. LA. a modeling language for process engineering—I. The formal framework. *Comput. Chem. Eng.* 14 (8), 813–846.
- Stephanopoulos, G., Johnston, J., Kriticos, T., Lakshmanan, R., Mavrovouniotis, M., Siletti, C., 1987. DESIGN-KIT: An object-oriented environment for process engineering. *Comput. Chem. Eng.* 11 (6), 655–674.
- Sundaram, A., Ghosh, P., Caruthers, J.M., Venkatasubramanian, V., 2001. Design of fuel additives using neural networks and evolutionary algorithms. *AIChE J.* 47 (6), 1387–1406.
- Suresh, P., Hsu, S.-H., Akkisetty, P., Reklaitis, G.V., Venkatasubramanian, V., 2010a. OntoMODEL: ontological mathematical modeling knowledge management in pharmaceutical product development, 1: conceptual framework. *Ind. Eng. Chem. Res.* 49 (17), 7758–7767.
- Suresh, P., Hsu, S.-H., Reklaitis, G.V., Venkatasubramanian, V., 2010b. OntoMODEL: ontological mathematical modeling knowledge management in pharmaceutical product development, 2: applications. *Ind. Eng. Chem. Res.* 49 (17), 7768–7781.
- Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., Hashimoto, T.B., 2023. Stanford alpaca: An instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca.
- Taskiran, N.P., Chakraborty, A., Mann, V., Venkatasubramanian, V., 2024a. Accelerating drug discovery through the automatic population of a pharmaceutical ontology using knowledge graphs. In: 2024 AIChE Annual Meeting. AIChE.
- Taskiran, N.P., Tsai, C.-E., Chakraborty, A., Mann, V., Venkatasubramanian, V., 2024b. A knowledge-graph-based pharmaceutical engineering chatbot for drug discovery. In: 2024 AIChE Annual Meeting. AIChE.
- Team, T.M.R., 2024. Introducing DBRX: A new state-of-the-art open LLM.
- Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., et al., 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Templeton, A., Conerly, T., Marcus, J., Lindsey, J., Bricken, T., Chen, B., Pearce, A., Citro, C., Ameisen, E., Jones, A., Cunningham, H., Turner, N.L., McDougall, C., MacDiarmid, M., Freeman, C.D., Sumers, T.R., Rees, E., Batson, J., Jermyn, A., Carter, S., Olah, C., Henighan, T., 2024. Scaling monosemanticity: Extracting interpretable features from Claude 3 sonnet. *Transform. Circuits Thread*.
- Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., et al., 2022. Llama: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al., 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al., 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Trinh, T.H., Wu, Y., Le, Q.V., He, H., Luong, T., 2024. Solving olympiad geometry without human demonstrations. *Nature* 625 (7995), 476–482.
- Turing, A.M., 1950. Computing Machinery And Intelligence. *Mind* LIX (236), 433–460.
- Udrescu, S.-M., Tegmark, M., 2020. AI Feynman: A physics-inspired method for symbolic regression. *Sci. Adv.* 6 (16), eaay2631.
- Ungar, L., Venkatasubramanian, V., 1990. Advance reasoning architectures for expert systems. The CACHE Corporation, Austin, TX.
- Vaidyanathan, R., Venkatasubramanian, V., 1995. Digraph-based models for automated HAZOP analysis. *Reliab. Eng. Syst. Saf.* 50 (1), 33–49.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30.
- Vedam, H., Venkatasubramanian, V., 1999. PCA-SDG based process monitoring and fault diagnosis. *Control Eng. Pract.* 7 (7), 903–917.
- Venkatasubramanian, V., 1989. Knowledge-based Systems in Process Engineering: case Studies in Heuristic Classification. The CACHE Corporation, Austin, TX.
- Venkatasubramanian, V., 2019a. Artificial intelligence in materials science: the good, the bad, and the ugly. In: 2019 AIChE Annual Meeting. AIChE.
- Venkatasubramanian, V., 2019b. The promise of artificial intelligence in chemical engineering: Is it here, finally? *AIChE J.* 65 (2), 466–478.
- Venkatasubramanian, V., 2024. What's in an embedding? Would a rose be any embedding smell as sweet? *arXiv*.
- Venkatasubramanian, V., Mann, V., 2022. Artificial intelligence in reaction prediction and chemical synthesis. *Curr. Opin. Chem. Eng.* 36, 100749.
- Venkatasubramanian, V., Rich, S., 1988. An object-oriented two-tier architecture for integrating compiled and deep-level knowledge for process diagnosis. *Comput. Chem. Eng.* 12 (9–10), 903–921.
- Venkatasubramanian, V., Zhao, C., Joglekar, G., Jain, A., Hailemariam, L., Suresh, P., Akkisetty, P., Morris, K., Reklaitis, G.V., 2006. Ontological informatics infrastructure for pharmaceutical product development and manufacturing. *Comput. Chem. Eng.* 30 (10–12), 1482–1496.
- Viswanathan, S., Johnsson, C., Srinivasan, R., Venkatasubramanian, V., Årzen, K.E., 1998a. Automating operating procedure synthesis for batch processes: Part I. Knowledge representation and planning framework. *Comput. Chem. Eng.* 22 (11), 1673–1685.
- Viswanathan, S., Johnsson, C., Srinivasan, R., Venkatasubramanian, V., Årzen, K.E., 1998b. Automating operating procedure synthesis for batch processes: Part II. Implementation and application. *Comput. Chem. Eng.* 22 (11), 1687–1698.
- Viswanathan, S., Shah, N., Venkatasubramanian, V., 2002. Hybrid framework for hazard identification and assessment in batch processes. *AIChE J.* 48 (8), 1765–1774.
- Viterbi, A., 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. Inf. Theory* 13 (2), 260–269.
- Von Stosch, M., Oliveira, R., Peres, J., de Azevedo, S.F., 2014. Hybrid semi-parametric modeling in process systems engineering: Past, present and future. *Comput. Chem. Eng.* 60, 86–101.
- Wei, J., Wei, J., Tay, Y., Tran, D., Webson, A., Lu, Y., Chen, X., Liu, H., Huang, D., Zhou, D., et al., 2023. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846*.
- Weininger, D., 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* 28 (1), 31–36.
- Weininger, D., Weininger, A., Weininger, J.L., 1989. SMILES. 2. Algorithm for generation of unique SMILES notation. *J. Chem. Inf. Comput. Sci.* 29 (2), 97–101.
- Weizenbaum, J., 1966. ELIZA—a computer program for the study of natural language communication between man and machine. *Commun. ACM* 9 (1), 36–45.
- Wilson, Z.T., Sahinidis, N.V., 2017. The ALAMO approach to machine learning. *Comput. Chem. Eng.* 106, 785–795.
- Winograd, T., 1971. Procedures as a representation for data in a computer program for understanding natural language. Technical Report, Massachusetts Inst of Tech Cambridge Project Mac.
- Zhang, Q., Ding, K., Lyv, T., Wang, X., Yin, Q., Zhang, Y., Yu, J., Wang, Y., Li, X., Xiang, Z., et al., 2024. Scientific large language models: A survey on biological & chemical domains. *arXiv preprint arXiv:2401.14656*.
- Zhang, H., Song, H., Li, S., Zhou, M., Song, D., 2022. A survey of controllable text generation using transformer-based pre-trained language models. *arXiv preprint arXiv:2201.05337*.
- Zhao, C., Bhushan, M., Venkatasubramanian, V., 2005a. PHASuite: an automated HAZOP analysis tool for chemical processes: part I: knowledge engineering framework. *Process Saf. Environ. Prot.* 83 (6), 509–532.
- Zhao, C., Bhushan, M., Venkatasubramanian, V., 2005b. PHASuite: An automated HAZOP analysis tool for chemical processes: Part II: Implementation and Case Study. *Process Saf. Environ. Prot.* 83 (6), 533–548.