

A multi-scale cross-domain thermochemical knowledge-graph

Sebastian Mosbach^{1,2}, Angiras Menon¹, Feroz Farazi¹, Nenad Krdzavac¹,
Xiaochi Zhou¹, Jethro Akroyd^{1,2}, Markus Kraft^{1,2,3}

released: 11 July 2020

¹ Department of Chemical Engineering
and Biotechnology
University of Cambridge
Philippa Fawcett Drive
Cambridge, CB3 0AS
United Kingdom

E-mail: mk306@cam.ac.uk

² CARES
Cambridge Centre for Advanced
Research and Education in Singapore
1 Create Way
CREATE Tower, #05-05
Singapore, 138602

³ School of Chemical
and Biomedical Engineering
Nanyang Technological University
62 Nanyang Drive
Singapore, 637459

Preprint No. 257



Keywords: Knowledge-graphs, quantum chemistry, thermodata, automation, interoperability

Edited by

Computational Modelling Group
Department of Chemical Engineering and Biotechnology
University of Cambridge
Philippa Fawcett Drive
Cambridge CB3 0AS
United Kingdom

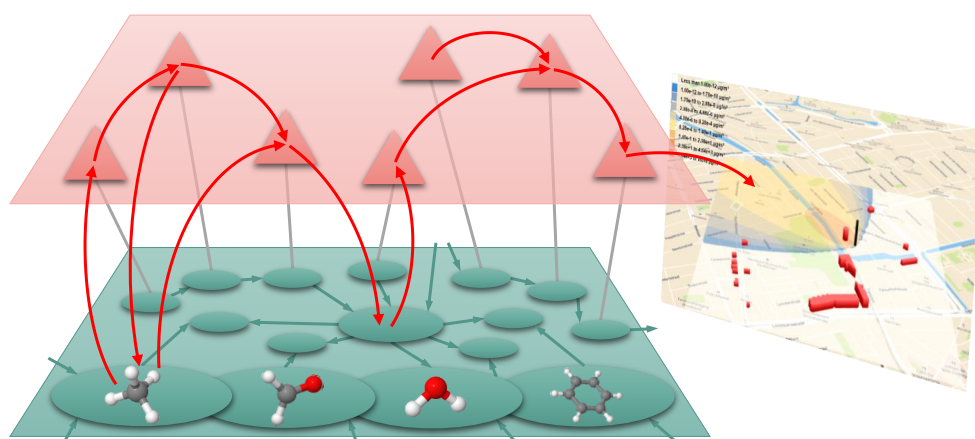
E-Mail: c4e@cam.ac.uk

World Wide Web: <https://como.ceb.cam.ac.uk/>



Abstract

In this paper, we develop a set of software agents which improve a knowledge-graph containing thermodynamic data of chemical species by means of quantum chemistry calculations and error-cancelling balanced reactions. The knowledge-graph represents species-associated information by making use of the principles of linked data as employed in the Semantic Web, where concepts correspond to vertices and relationships between the concepts correspond to edges of the graph. We implement this representation by means of ontologies, which formalise the definition of concepts and their relationships, as a critical step to achieve interoperability between heterogeneous data formats as well as software. The agents, which conduct quantum chemistry calculations and derive estimates of standard enthalpies of formation, update the knowledge-graph with newly obtained results, improving data values and adding nodes and connections between them. A key distinguishing feature of our approach is that it extends an existing, general-purpose knowledge-graph, called J-Park Simulator (theworldavatar.com), and its eco-system of autonomous agents, thus enabling seamless cross-domain applications in wider contexts. To this end, we demonstrate how quantum calculations can directly affect the atmospheric dispersion of pollutants in an industrial emissions use-case.



Highlights

- Autonomous agents for quantum chemistry and enthalpy of formation calculations are developed.
- Enthalpies of formation are derived using error-cancelling balanced reactions.
- The agents are integrated, exploiting knowledge-graph-enabled interoperability.
- A multi-scale, cross-domain industrial pollution use-case is considered.
- Quantum calculations are directly seen to affect atmospheric pollutant dispersion.

Contents

1	Introduction	3
2	The World Avatar	4
3	Design and implementation	6
3.1	The chemical knowledge-graph	6
3.2	Agent template	7
3.3	Error-cancelling balanced reaction agent	9
3.4	Quantum calculation agent	11
3.5	Thermodata agent	12
3.6	Integration	14
3.7	Monitoring agent	14
4	Use-case: Atmospheric pollutant dispersion	16
5	Conclusions	17
	References	19

1 Introduction

Optimising industrial operations is critical for mitigating their environmental impact, that is, maximising energy efficiency and minimising pollution and waste of resources. In order for this to be most effective, systems must not be considered in isolation, but rather connected to one another as part of a network. Furthermore, implementing this requires models of all systems involved. Such reasoning naturally leads to the ideas of Industry 4.0 [29] and the Internet of Things (IoT) [2], where digital twins of real objects can communicate with each other via the internet.

In case one or more of the models in the network involve chemistry in some form, chemical models need to be built, using chemical data about species and reactions. Numerous databases exist for such purposes, containing various levels of detail. For instance, some of the most widely used are Reaxys [22], PubChem [26], and the CAS Registry [1], each of which include a wealth of chemical and physical substance information on in excess of 10^8 compounds. Amongst the largest is the Chemical Universe Database GDB-17 [41] with more than 10^{11} structures, though it contains SMILES strings [46] only. At a smaller scale, PrIme [17] is largely focussed on combustion, and includes not only a warehouse of experimental and computational data, but also an associated set of tools for a variety of tasks related to model development. At the most detailed end of the spectrum, and even more focussed on thermodynamic properties, are collections that include quantum chemistry calculations. For example, Ramakrishnan et al. [40] have conducted quantum calculations for a GDB-17 subset of more than 10^5 structures. NIST’s CCCBDB [24] provides extensive experimental and quantum calculation data on the thermochemistry of 1968 small gas-phase species. Active Thermochemical Tables (ATcT) [42] use a statistical approach to synthesise accurate and consistent thermodata from experiments and computations for 1617 species. Highly accurate thermochemistry has been calculated for 219 molecules relevant to combustion [21].

With so much data available from so many different sources, inconsistencies are ubiquitous, both in terms of naming of species and in terms of data values of thermodynamic, transport, and kinetic properties of species and reactions (see *e.g.* [28]). Combined with the complexity of fuel combustion models, with hundreds or even thousands of species and reactions, it is clear that automation is inevitable. A number of attempts have been made to automate the generation of kinetic mechanisms, and in particular the generation of thermodynamic data. The Reaction Mechanism Generator (RMG) [20] constructs species and reactions based on a known set of rules, which also includes estimation of thermodynamic, kinetic, and transport properties. Building upon RMG, Keçeli et al. [25] have developed a Predictive Automated Computational Thermochemistry (PACT) software package which automatically generates thermodynamic data from quantum calculations for species involved in the combustion of an arbitrary hydrocarbon fuel. Li et al. [31] have implemented a self-evolving thermochemistry machine which uses a convolutional neural network to predict species formation enthalpies. The network is trained on a database of density functional theory (DFT) calculations and estimates its own prediction errors, launches new DFT calculations whenever the error exceeds a given threshold, and includes the new results into its training database, thus continuously and automatically improving its predictions.

As the above examples show, automating specific, well-defined tasks can be achieved through tailor-made code. However, the challenge becomes much harder as soon as one considers more general contexts, where, in order to make substantial progress with automation, a more fundamental, systematic approach is required. One such approach utilises knowledge-graphs, which represent information by making use of the principles of Linked Data [3, 5] as employed in the Semantic Web [4], where concepts correspond to vertices and relationships between concepts correspond to edges of the graph. This representation is implemented by means of ontologies, which formalise the definition of concepts and their relationships through collections of subject-predicate-object triples. The power of knowledge-graphs at a large scale has become abundantly clear in various applications over the last few years [16, 36]. In the field of chemistry, it has been recognised that subject-predicate-object representations of chemical data are of value, in particular for automation [43], with PubChemRDF [19] being one of many examples.

In previous work on knowledge-graphs in the chemistry domain [34], we have developed an ontology for quantum chemistry calculations based on Chemical Markup Language (CML [39]), called OntoCompChem [27]. We have furthermore developed an ontology for chemical reaction mechanisms, called OntoKin [13]. In order to connect these two worlds, linking reaction mechanisms and quantum chemistry [14], we have introduced an ontology for unique chemical species, called OntoSpecies, thus integrating the union of all these concepts into a single knowledge-graph.

The purpose of the present paper is to employ a knowledge-graph approach, specifically with newly-developed software agents, to improve thermodynamic data for chemical species, and apply it to a multi-scale, multi-domain combustion example. Whilst this process can also be automated by means of other methods, our approach achieves this in a fashion that, by design, guarantees interoperability between data and software, and thus allows seamless integration with complex cross-domain applications. As such, the knowledge-graph approach renders automated thermochemistry extensible beyond its original scope.

The paper is structured as follows. Section 2 explains the context, namely the general-purpose knowledge-graph within which the present work is placed. In section 3, we outline a generic agent design and describe a number of thermochemical agents and how they are integrated. In section 4, we provide a cross-domain use-case of atmospheric dispersion of industrial pollution. Conclusions are drawn in section 5.

2 The World Avatar

The term ‘World Avatar’ intends to capture the idea of representing every aspect of the real world in a digital ‘mirror’ world. This is essentially an extension of the Digital Twin concept, where, taking an example from Industry 4.0, a device or a unit operation in an industrial process has a corresponding virtual representation. Lifting the restriction to industrial devices, thus considering virtualisation of any abstract concept or process, is a logical continuation, similar to extending the Internet of Things to the Internet of Services and beyond.




Figure 1: Current design of the J-Park Simulator (JPS) as an implementation of a World Avatar. Agents are part of the knowledge-graph and operate on it.

The J-Park Simulator (JPS)¹ [11] is an implementation of the World Avatar concept. Figure 1 illustrates its main underlying principles. At the heart of the JPS lies a knowledge-graph that is intended to be general-purpose and all-encompassing. As a representation of data, the key distinguishing feature of a knowledge-graph is that individual items of information are linked to each other. In the JPS, this Linked Data is achieved through the use of Internationalised Resource Identifiers (IRIs), essentially generalised web-addresses, in line with the Semantic Web. The concepts in the knowledge-graph and the links between them are implemented by means of ontologies for various domains. These include process engineering (OntoCAPE [32]), Eco-Industrial Parks (EIPs) (OntoEIP [47–49]), electrical power systems (OntoPowerSys [10]), and 3D buildings and landscapes (OntoCityGML [11]). In the chemistry domain, we have developed ontologies for the sub-domains of quantum chemistry (OntoCompChem [27]), species (OntoSpecies [14]), and kinetic reaction mechanisms (OntoKin [13]). In addition, various subgraphs of the Linked Open Data (LOD) Cloud² are also connected to the JPS knowledge-graph, in particular DBpedia³ [30].

Beyond mere data representation, the JPS contains an eco-system of software agents that act autonomously and operate on the knowledge-graph, constantly updating it, as also illustrated in Fig. 1. Crucially, the agents themselves are part of the knowledge-graph, governed by the OntoAgent [50] ontology. As agents are continuously operating on the knowledge-graph, it evolves in time. In particular, we have developed agents for automatic agent discovery and composition [50], *i.e.* agents that create new, composite agents for more complex tasks. And further, in order to facilitate the usage of agents and simplify identification of an agent suitable for a specific task in an agent-rich environment, with

¹<http://www.theworldavatar.com/>

²<https://lod-cloud.net/>

³<https://wiki.dbpedia.org/>

an abundance of services available, we have established an agent market place based on block-chain technology and Smart Contracts [51].

The JPS started with a focus on virtualising industrial operations within the Jurong Island EIP in Singapore [37, 38, 48], but has since expanded well beyond this original scope. It has been employed in a variety of cross-domain applications that require interoperability between heterogeneous software and data formats, such as optimal site selection of nuclear power plants [9] and industrial pollution prediction of ships and power plants [11]. In addition, scenario planning is possible using the Parallel World Framework [12], which generically allows asking what-if questions and exploring alternatives in complex multi-domain applications.

3 Design and implementation

In this section, we describe the agents that form the basis of the present paper. We give some theoretical background of what they do and outline their architecture as well as some key features of their implementation. We also explain how they are integrated, *i.e.* how they work with each other and the knowledge-graph, and thus how they are part of the wider World Avatar agent eco-system as described in the previous section.

3.1 The chemical knowledge-graph

In order to provide the necessary context of what environment the agents are operating in, we briefly recall here what chemical entities are represented in the knowledge-graph and how. As mentioned above, we have previously created three ontologies for a number of chemical concepts, populated triple-stores with collections of instances, and established links between them:

The OntoKin [13] ontology covers chemical reaction mechanisms – collections of chemical elements, species, and reactions together with their associated thermodynamic, kinetic, and transport data. A web-interface is available⁴ that allows up- and download, and basic queries [15]. In addition, we have developed an API that allows programmatic up- and downloading, and import of, export of, and conversion between widely used mechanism file formats.

The OntoCompChem [27] ontology is concerned with representing quantum calculations. Whilst the ontology itself is not intrinsically specific to a particular piece of software, the parser and API we have written for file import is designed for Gaussian 09/16 calculations. A web-interface is also available⁵. Representable quantities include empirical formula, InChI and SMILES strings, molecular geometry, level of theory and basis set, rotational and vibrational frequencies, rotational constants, and electronic energies.

Thirdly, we have introduced the OntoSpecies ontology [14], primarily for the purpose of identifying chemical species uniquely. The ontology captures basic physical and chemical

⁴<http://www.theworldavatar.com/ontokin/>

⁵<http://www.theworldavatar.com/molhub/>

information about species, such as elemental composition, connectivity between atoms, molecular geometry, and in particular standard enthalpy of formation. In itself, OntoSpecies is not meant to and does not need to store much information – its main value lies in providing links via IRIs, connecting quantum calculations with each other as well as species in mechanisms, which are labelled using arbitrary strings. These IRI links are used to disambiguate arbitrary chemical labels, distinguishing different species with identical labels as well as recognising identical species with different labels, thus circumventing naming inconsistencies.

The agents described in the following critically depend upon this infrastructure, with its links between relevant concepts, being in place.

3.2 Agent template

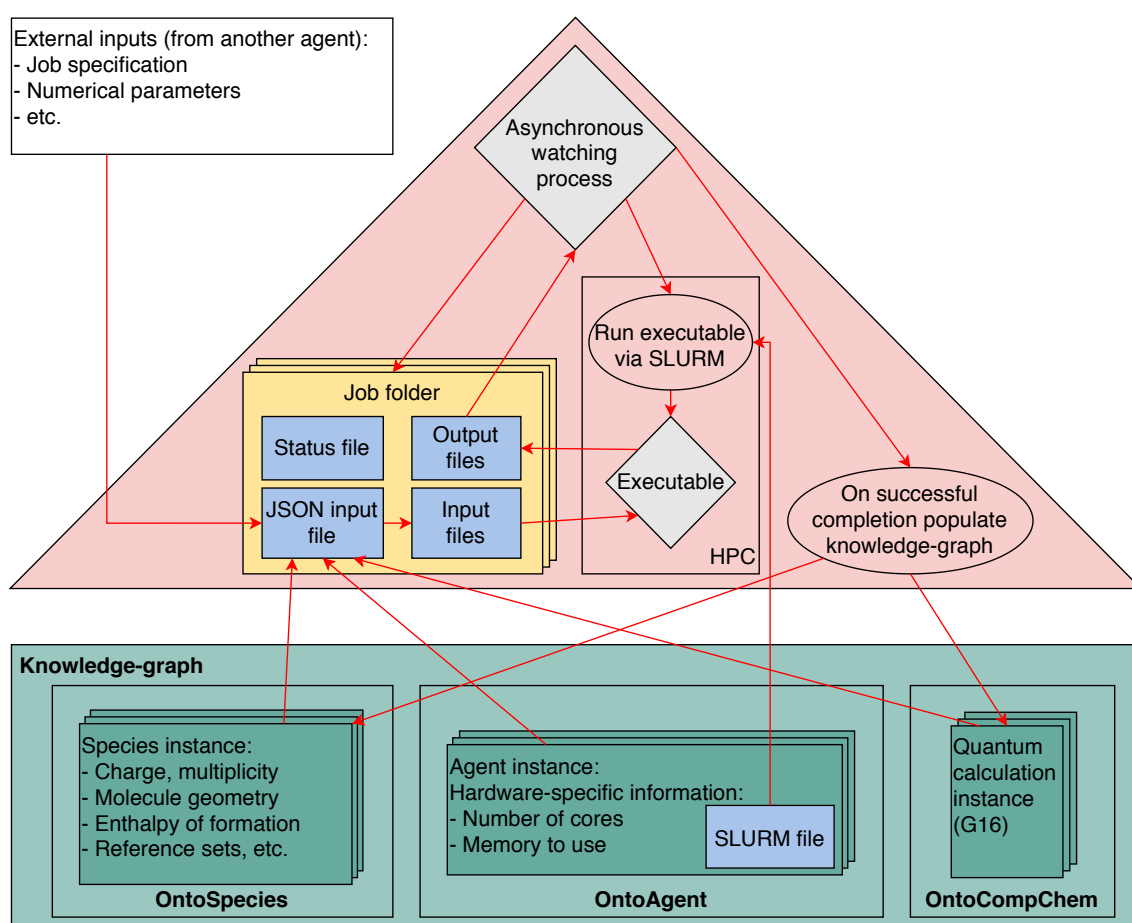


Figure 2: Elements of a generic agent (red triangle) and how they interact with the knowledge-graph (green box). An asynchronous watcher (grey diamond) manages running an executable (grey diamond), with all associated input and output files (blue boxes).

When developing agents for various tasks, one quickly arrives at the realisation that many design features are entirely generic, *i.e.* independent of the specific task. Such design

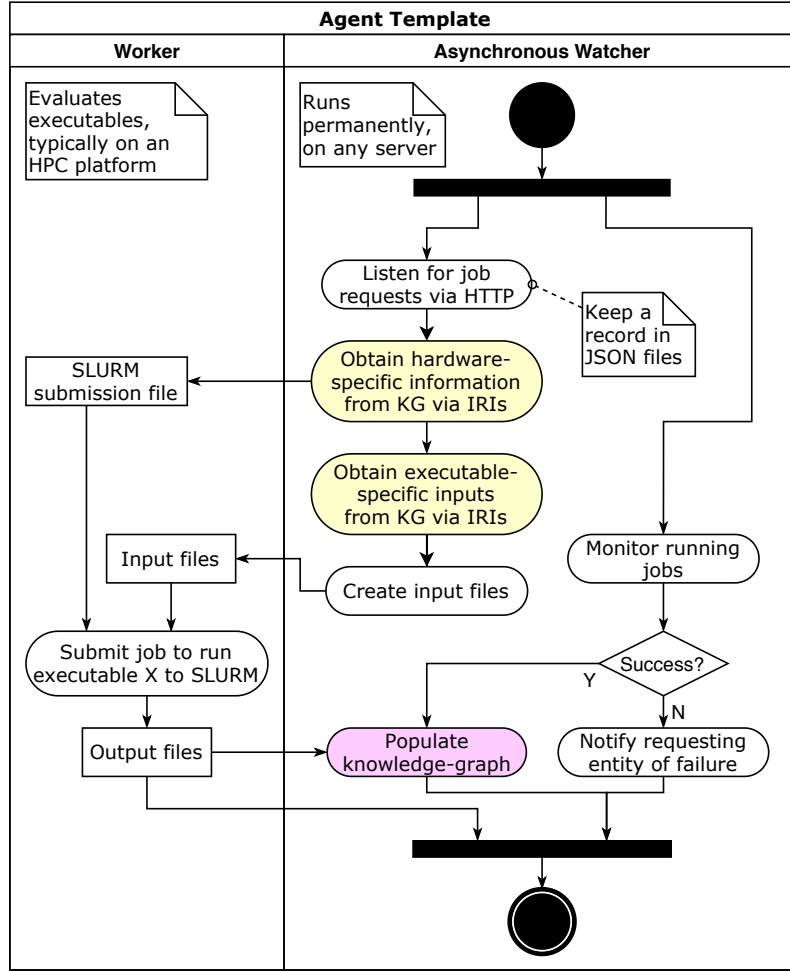


Figure 3: UML activity diagram of a generic agent which enables computational jobs (running of an executable “X”) to be executed asynchronously on an HPC system upon HTTP request, with inputs obtained from the knowledge-graph (shaded in yellow) and outputs written back into the knowledge-graph (shaded in magenta).

features include most importantly listening and responding to HTTP requests, submission and monitoring of jobs to a resource manager (e.g. SLURM) on an HPC system, and managing input and output files associated to a job as well as their transfer between the hardware platforms involved. We have thus developed a template that is, within reason, applicable to ‘any’ (Linux) executable.

Figure 2 illustrates the main components of this generic agent and how it interacts with the knowledge-graph. The agent (the triangle shaded in red) consists of two executable elements: An asynchronous watching process and a generic executable. The asynchronous watcher handles HTTP requests and responses, through which other agents can request jobs. Secondly, it manages input and output files, which are being stored in separate folders associated to each job. And thirdly, it takes care of submitting jobs to a resource manager tool on an HPC system as well as monitoring any jobs that are running. It does

this using a status file associated to each job in its corresponding folder. This design is robust to unexpected interruptions such as power-cuts or reboots by a system administrator. Once the hardware and the agent are restarted, the latter will continue its operations unaffected.

When the agent receives job request via HTTP, inputs to the agent are passed to it in Semantic Web standard JavaScript Object Notation (JSON) format. These inputs can include ‘direct’ inputs (from other agents) such as parameters specifying the task to be executed, but more typically will take the form of IRIs, pointing to relevant entities that are represented in the knowledge-graph. Also retrieved from knowledge-graph are parameters that are specific to the HPC hardware on which the executable that constitutes the agent is to be executed, such as the number of CPU-cores and amount of memory to be used, as well as a SLURM file which is used for submission to the resource manager.

As an aside, we note that this design implies that only a single instance of this agent needs to be deployed on a particular piece of hardware, and also in general, if the volume of incoming requests is sufficiently small.

Figure 3 shows a Unified Modelling Language (UML) activity diagram, illustrating the step-by-step activities of the components. Handling HTTP requests and monitoring running jobs are conducted asynchronously in parallel, as any events can happen at any time, in any order. Most agents will retrieve inputs from the knowledge-graph (shaded in yellow) and write outputs back to the knowledge-graph (shaded in magenta). Which inputs and outputs exactly are exchanged with the knowledge-graph is of course not generic, but depends on the executable in question. Another critical activity relates to job monitoring: In practice, there are many things that can go wrong in the submission and execution of a computational job. It is therefore important to have robust procedures in place to deal with any unexpected failure in a controlled fashion, which in particular involves escalating useful error messages back to the calling entity, and ultimately a human, if the cause of the problem cannot be rectified by means of an algorithm.

3.3 Error-cancelling balanced reaction agent

The purpose of this agent is two-fold. Given a set of chemical species, it can produce estimates of standard enthalpy of formation, and in addition, within the given set, identify individual species as either consistent or inconsistent. The agent operates in two steps. The first step computes the standard enthalpy of formation of a species by using error-cancelling balanced reactions (EBRs) [6]. At the centre of this method lies Hess’ law, which states that the total enthalpy change of a reaction, ΔH_r° is equal to the sum of all individual enthalpy changes, independent of the reaction pathway:

$$\Delta H_r^\circ = \sum_{s \in S^P} \nu(s) \Delta H_f^\circ(s) - \sum_{s \in S^R} \nu(s) \Delta H_f^\circ(s) \quad (1)$$

Here, S^P and S^R denote the sets of products and reactants, $\nu(s)$ the stoichiometric coefficients, and $\Delta H_f^\circ(s)$ the standard enthalpy of formation of species s . The reaction enthalpy can be estimated from the results of quantum chemistry calculations via

$$\Delta H_r^\circ = \sum_{s \in S^P} \nu(s) E^\circ(s) - \sum_{s \in S^R} \nu(s) E^\circ(s), \quad (2)$$

where $E^\circ(s)$ is the zero-point corrected ground state energy given by quantum chemical calculations. If the enthalpy of reaction is computed as above, then the standard enthalpy of formation of a given target species, s_T can be estimated by rearranging Hess’ law:

$$\Delta H_f^\circ(s_T) = \frac{1}{v(s_T)} \sum_{s \in S^p} v(s) \Delta H_f^\circ(s) - \frac{1}{v(s_T)} \left(\Delta H_r^\circ + \sum_{s \in S^R \setminus \{s_T\}} v(s) \Delta H_f^\circ(s) \right) \quad (3)$$

However, it is known that estimating enthalpies using quantum chemical calculations results in systematic errors [45]. To circumvent this, error-cancelling balanced reactions (EBRs), are employed, which make use of structural and electronic similarities between the species in a reaction to allow cancellation of the systematic errors. Examples of types of EBRs include isogyric reactions, which conserve the number of spin states during the reaction, and isodesmic reactions, which conserves the number of each type of bond during the reaction. Using EBRs enables improved estimates of standard enthalpies of formation of species from quantum chemical calculations.

The second step performed by the agent is a heuristic cross-validation. The algorithm is described in detail elsewhere [6], and so only a brief summary is provided here. A set of species is given to the agent, each with a reference value for the standard enthalpy of formation. These reference values may be derived experimentally or from high-level computational methods. Next, one species in the set is selected, in a leave-one-out cross-validation method. For the selected species, a user-defined number of EBRs are generated. Each EBR can then be used to estimate the standard enthalpy of formation of the species by using Eqn. (3), an estimate of ΔH_r° derived from quantum calculations, and the reference enthalpies for all other species in the EBR. For each EBR, the estimated standard enthalpy of formation generated for the species is then compared to the reference value, and an error is computed:

$$\epsilon_r(r, s_T) = \left| \Delta H_{f, \text{reference}}^\circ(s_T) - \Delta H_f^\circ(r, s_T) \right| \quad (4)$$

The EBR is then accepted if this error is below a defined upper limit, ϵ_r^{\max} , and rejected otherwise. If all EBRs for a given species are rejected, the species is flagged as potentially inconsistent. This process is repeated for all of the species in the reference set to sort them as consistent or potentially inconsistent.

The set of consistent and potentially inconsistent species is then refined by selecting the potentially inconsistent species with the highest average error and generating a new set of EBRs for it using only species that were deemed consistent in the initial sorting. The errors are then recomputed for this new set of EBRs. If the average error for the new set of EBRs, $\bar{\epsilon}_r^{\text{new}}$, is lower than the error for the initial set of EBRs, $\bar{\epsilon}_r^{\text{initial}}$, then it is assumed that the original inconsistency for this species is due to another species that appeared in the original EBRs, and the selected species is added to the consistent set. The process is repeated for each potentially inconsistent species in descending order of average error until the set of inconsistent species does not change. At this point, convergence is reached, and the set of inconsistent species is reported as needing an improved estimate of the standard enthalpy of formation. It is noted that, depending on the size of the species set and other parameters, executing the algorithm can be computationally expensive.

The design and implementation of this EBR agent follows that of the generic template outlined in subsection 3.2, as illustrated in Figs. 2 and 3. Jobs are requested via HTTP.

The JSON-formatted inputs, that are passed as arguments as part of the HTTP request, include a list of pairs of IRIs. Each pair consists of an IRI to a unique species instance in OntoSpecies, and an IRI to a quantum calculation in OntoCompChem. This list defines the species set that the EBR algorithm is applied to. All associated information, like reference values for the standard enthalpies of formation for each species, can be retrieved from the knowledge-graph using the provided IRIs. Other JSON inputs include for example the type of EBR to be used (isogyric, isodesmic, *etc.*), and algorithm termination criteria (maximum numbers of iterations of various loops). Once all required information has been assembled, input files for the core EBR code are generated and the job is submitted to the resource manager on an HPC platform. In case of successful termination, the resulting new estimates of standard enthalpies of formation as well as the finding of whether or not a species is being deemed consistent or not are stored in the knowledge-graph (as part of the relevant instances in OntoSpecies).

3.4 Quantum calculation agent

The purpose of this agent is to conduct a quantum chemistry calculation using the Gaussian [18] software. Quantum chemical calculations are used to derive the molecular properties of a chemical system from first principles by solving the time-independent Schrödinger Equation. Solving this equation for a molecule, *i.e.* for a system consisting of several electrons and nuclei, yields the wavefunction of the system, which provides information about the quantum state of the system, such as the positions of nuclei and electrons and the energy associated with their particular configuration. The Schrödinger Equation can be written as

$$\hat{H}\Psi(\vec{r},\vec{R}) = E\Psi(\vec{r},\vec{R}), \quad (5)$$

where \hat{H} is the Hamiltonian operator, Ψ is the wavefunction of the system, \vec{r} and \vec{R} are the positions of electrons and nuclei, and E is the eigenvalue, representing the total energy of the system. The Hamiltonian operator can be written as

$$\hat{H} = \hat{T}_n + \hat{T}_e + \hat{V}_{n-n} + \hat{V}_{n-e} + \hat{V}_{e-e}, \quad (6)$$

where, \hat{T}_n and \hat{T}_e operators represent the kinetic energy of the nuclei and electrons respectively, and \hat{V}_{n-n} , \hat{V}_{n-e} and \hat{V}_{e-e} represent the potential energy of the nucleus-nucleus, nucleus-electron, and electron-electron interactions, respectively.

Typically, the Born-Oppenheimer approximation is invoked, which neglects the coupling between the electrons and nuclei of a system. This results in the electronic Schrödinger equation, which is what modern computational chemistry packages like Gaussian solve to derive molecular properties:

$$\hat{H}_e\Psi(\vec{r},\vec{R})_e = E_e\Psi(\vec{r},\vec{R})_e \quad (7)$$

In this notation, $\Psi(\vec{r},\vec{R})_e$ is the electronic wavefunction, E_e are the eigenvalues representing the electronic energies of the system, and \hat{H}_e is the electronic Hamiltonian:

$$\hat{H}_e = \hat{T}_e + \hat{V}_{n-e} + \hat{V}_{e-e} \quad (8)$$

For chemical systems larger than simple hydrogen-like atoms, the Schrödinger equation can only be solved numerically. Various methods to do so are implemented in modern computational chemistry software such as Gaussian. One very popular method is Density Functional Theory (DFT), which derives the properties of a chemical system by manipulating the electronic Schrödinger equation and solving for the energy as a function of the electron density $\rho(\vec{r})$ of the system [52]. This results in the equation [23]

$$E[\rho(\vec{r})] = T_s[\rho(\vec{r})] + V_{n-e}[\rho(\vec{r})] + J[\rho(\vec{r})] + E_{xc}[\rho(\vec{r})], \quad (9)$$

where $T_s[\rho(\vec{r})]$ is the kinetic energy of the non-interacting model system and $J[\rho(\vec{r})]$ is the Coulomb energy functional, both of which are known exactly. Similarly, $V_{n-e}[\rho(\vec{r})]$ is the nucleus-electron attraction potential functional whose form can be derived given the Born-Oppenheimer approximation. The final term, E_{xc} is termed the exchange-correlation functional, and accounts for the difference in kinetic and potential energies between the real interacting-electron system, and the approximate non-interacting electron system.

DFT is a very popular approach, as the electron density is only a function of the spatial coordinates, making it a computationally efficient way to derive molecular properties through quantum chemistry, and is hence adopted in this work. Nevertheless, such methods still require substantial computational resources, and do typically need to be run on HPC systems.

The quantum calculation agent automates the solutions of the above equations via Gaussian. Similar to the enthalpy of formation agent, the implementation of this agent follows exactly the design outlined in subsection 3.2 (Figs. 2 and 3), *i.e.* calculations are requested via HTTP and are conducted on an HPC system, with key inputs and outputs retrieved from and written back into the knowledge-graph. JSON inputs passed to this agent include in particular an IRI of a unique species instance in OntoSpecies. Species definition information required for a quantum calculation such as charge, spin multiplicity, and geometry of the molecule are retrieved from the knowledge-graph by means of SPARQL queries. Having obtained all necessary inputs from the knowledge-graph, a plain-text Gaussian input file is then populated. Job submission and monitoring are managed generically by an asynchronous watching process. In case of successful completion, the result of the quantum calculation, namely the log file output by Gaussian, is converted to OWL and a new instance of the G16 class, as defined in the OntoCompChem ontology, is created in the knowledge-graph.

3.5 Thermodata agent

The purpose of this agent is to calculate three thermodynamic quantities for a chosen chemical species as functions of temperature T , namely the heat capacity at constant pressure C_p , enthalpy H , and entropy S [14]. These thermodynamic quantities are derived by means of the molecular partition function, $q = q_T q_V q_R q_E$, whose components consist of the translational (q_T), vibrational (q_V), rotational (q_R), and electronic (q_E) partition functions derived from standard statistical mechanics expressions and the rigid-rotor-

harmonic-oscillator (RRHO) approximation [33]:

$$q_T = \left(\frac{mk_B T}{2\pi\hbar^2} \right)^{\frac{3}{2}} V \quad (10)$$

$$q_V = \prod_{i=1}^{N_V} \frac{\exp\left(-\frac{h\nu_i}{2k_B T}\right)}{1 - \exp\left(-\frac{h\nu_i}{k_B T}\right)} \quad (11)$$

$$q_R = \frac{(8\pi^3 I_x I_y I_z)^{1/2} (k_B T)^{3/2}}{\sigma \pi \hbar^3} \quad (12)$$

$$q_E \approx g_0^E \quad (13)$$

Here, h denotes Planck's constant, k_B is the Boltzmann constant, m is the mass, V is the volume, ν_i is the magnitude of the i^{th} frequency, I_k is the moment of inertia about the k^{th} axis, σ is the symmetry number, and g_0^E is the ground state electronic degeneracy. The necessary frequencies, rotations, and ground state energies to compute the partition functions are obtained from the results of quantum chemistry calculations. Once the molecular partition function is constructed, the molar heat capacity, entropy, and enthalpy can be derived from the following expressions:

$$C_p = C_v + N_A k_B \quad \text{with} \quad C_v = N_A k_B T \frac{\partial^2(T \ln q)}{\partial T^2} \quad (14)$$

$$S = N_A k_B \left[\frac{\partial(T \ln q)}{\partial T} - \ln N_A + 1 \right] \quad (15)$$

$$\Delta H = \int_0^T C_p dT = \frac{N_A k_B T^2}{q} \frac{\partial q}{\partial T} + N_A k_B T \quad (16)$$

Here, N_A denotes Avogadro's number. Note that Eqn. (16) provides instead of an absolute value only an enthalpy difference. In order to obtain meaningful absolute values, it is therefore necessary to provide a reference value for the enthalpy of formation at a known reference temperature, usually 298.15 K. In practice, the functional dependence of the three thermodynamic quantities on temperature is captured by fitting the widely-used NASA polynomials to the thermochemical data extracted from the partition functions, with seven coefficients c_1, \dots, c_7 for each polynomial, for two contiguous temperature ranges:

$$\frac{C_p}{R} = c_1 + c_2 T + c_3 T^2 + c_4 T^3 + c_5 T^4 \quad (17)$$

$$\frac{H}{R} = c_1 T + \frac{c_2}{2} T^2 + \frac{c_3}{3} T^3 + \frac{c_4}{4} T^4 + \frac{c_5}{5} T^5 + c_6 \quad (18)$$

$$\frac{S}{R} = c_1 \ln T + c_2 T + \frac{c_3}{2} T^2 + \frac{c_4}{3} T^3 + \frac{c_5}{4} T^4 + c_7 \quad (19)$$

The calculations involved for the thermodata agent are computationally very light and thus do not require submission to an HPC system. Instead, they can be conducted directly by whichever webservice responds to the HTTP request. Therefore, this agent can follow a simpler design than the one outlined in subsection 3.2. When it receives an HTTP request, the JSON arguments need to contain an IRI of a quantum calculation in the knowledge-graph from which thermodata is to be derived, as well as an IRI to a unique species instance in order to retrieve a reference value for the enthalpy of formation.

3.6 Integration

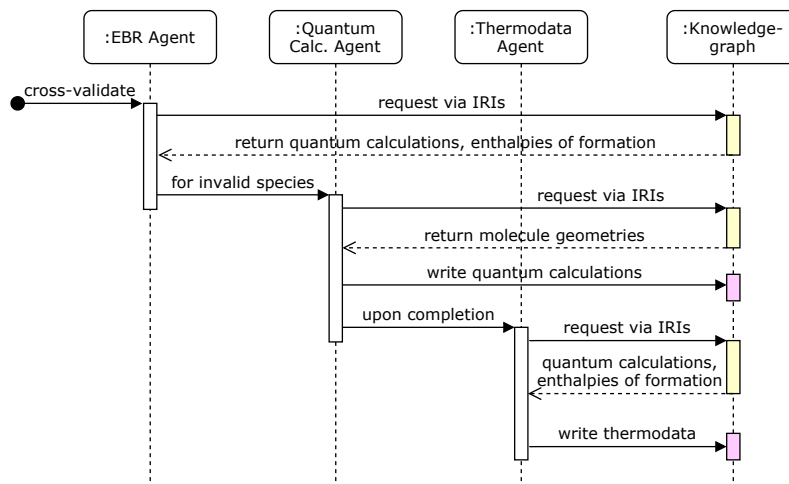


Figure 4: UML sequence diagram illustrating the interaction of the agents and how they act on the knowledge-graph (inputs to agents from knowledge-graph shaded in yellow, outputs back into the knowledge-graph shaded in magenta).

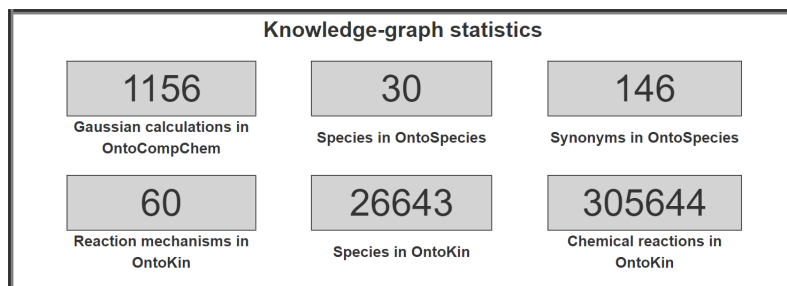
Figure 4 sketches in a UML sequence diagram how the three agents described above are being integrated, *i.e.* how they communicate with each other and how they interact with the knowledge-graph. The basic idea is a three-step process: Firstly, the EBR agent commences the process by retrieving previously obtained quantum calculation results and enthalpies of formation for a given set of species from the knowledge-graph via IRIs, and then detecting inconsistent species amongst the given set using cross-validation. Secondly, for each of those species that have been identified as inconsistent, the EBR agent requests quantum calculations from the quantum calculation agent at a higher level of theory than currently available in the knowledge-graph. Lastly, for any quantum calculation job that has finished successfully, the thermodata agent derives the corresponding thermodata. This involves retrieving a consistent estimate of the standard enthalpy of formation from the knowledge-graph, as populated by the EBR agent. As the final step, the thermodata agent then updates a chemical mechanism if it contains a species that is linked to the unique species whose thermodata was just updated.

As mentioned in subsection 3.1, we note that this sequence of three agents reading from and writing to the knowledge-graph is made possible by and relies upon IRI links being present in the knowledge-graph between quantum calculations (OntoCompChem), unique species (OntoSpecies), and species as they are part of chemical mechanisms (OntoKin).


3.7 Monitoring agent

In order to monitor the health status of the knowledge-graph, and to provide a high-level overview of how it evolves over time, we have created a web-page displaying various metrics⁶. Figure 5 displays screen-shots of the relevant parts of that web-page. The page

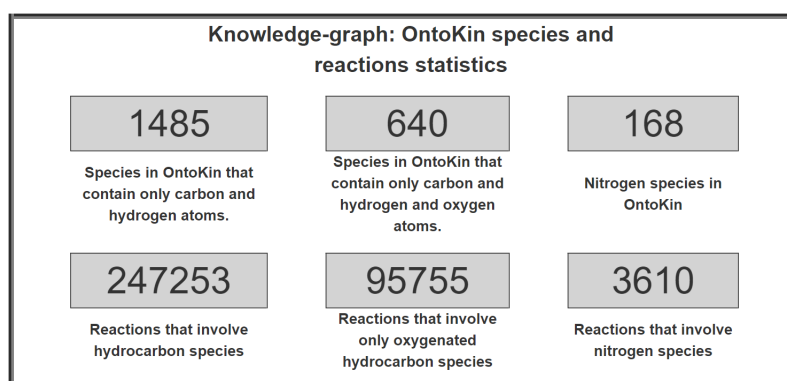
⁶<http://theworldavatar.com/graph/statisticsAction/>



(a) Counters of selected concepts, including quantum calculations, unique species, mechanisms, species, and reactions.



(b) Time-history chart showing numbers of species added on particular dates.




(c) Counters of selected types of species and reactions within chemical mechanisms (as represented in OntoKin).

Figure 5: Screen-shots of a web-page showing knowledge-graph statistics.

mainly features counters of instances of selected concepts in the OntoKin (mechanisms, species, and reactions), OntoSpecies (unique species), and OntoCompChem (quantum calculations) ontologies, as shown in Fig. 5(a), that can be useful for monitoring and diagnostics. A time-history chart of additional species within mechanisms (OntoKin) and as quantum calculations (OntoCompChem), showing how many instances were added on a particular date, is also included (Fig. 5(b)). More specialised statistics on certain types of species and reactions within mechanisms in OntoKin are shown in Fig. 5(c).

All the data displayed on this web-page is ultimately obtained through SPARQL queries from various triple-stores. In this context, it should be noted that subject-predicate-object triple representations of complex data structures such as chemical mechanisms involve considerable numbers of triples: For example, the sixty mechanisms we uploaded for demonstration purposes require several tens of millions of triples. The complexity of a query and the size of the repository both impact the response times. For this reason, in order to reduce loading times, most of the numbers shown on the page are cached. The cache is being refreshed once every 24 hours by this agent.

4 Use-case: Atmospheric pollutant dispersion




(a) CO emissions predicted using the original mechanism.



(b) CO emissions predicted using a mechanism with updated thermodata.



(c) uHC emissions predicted using the original mechanism.



(d) uHC emissions predicted using a mechanism with updated thermodata.

Figure 6: Screen-shots of a web-interface showing pollutant concentration distributions in the atmosphere over The Hague as predicted by ADMS, with 3D-representation of selected buildings (map data © Google).

As a use-case, we integrate the agents discussed in the previous sections into a cross-domain simulation of the atmospheric dispersion of pollutant emissions from a power plant and display this interactively on a web-page⁷. We consider a power plant that is fictitious, but inspired by the *Energiecentrale* in The Hague (see [11] for more details). Geographic data are processed using Google Maps API. Three-dimensional representations of selected buildings are included, which are stored in the knowledge-graph using

⁷<http://www.theworldavatar.com/JPS/>

the OntoCityGML ontology. The atmospheric dispersion of pollutants is simulated using the Atmospheric Dispersion Modelling System (ADMS) [7], which is based on a fluid-dynamic model that includes Gaussian plume air dispersion. ADMS requires as one of its inputs weather data, which are retrieved in real time from the world wide web and stored in the knowledge-graph by a dedicated weather agent. All relevant data are held in the knowledge-graph and as such are linked through IRIs.

Combustion in the stationary power generator is simulated by means of a Stochastic Reactor Model (SRM) [8] agent. The SRM has been developed primarily as an internal combustion engine (ICE) simulation tool and has been applied as such in numerous studies (see for example [35], and references therein). It uses detailed chemical kinetics and is thus able to predict gaseous and particulate pollutants. We consider a Primary Reference Fuel (PRF) as a surrogate fuel and employ as a mechanism a reduced PRF scheme that was specifically designed for ICE simulations [44]. The mechanism is represented in the knowledge-graph using OntoKin, is retrieved as an OWL file via IRI, and then converted into an SRM-readable format.

The process begins with the EBR agent cross-validating a set of 34 hydrocarbon species and concluding that, based on the given data using isodesmic reactions, 14 of the species are inconsistent. The EBR agent then launches quantum calculations for each of these species at a higher level of theory. Once a job finishes, the thermodata agent is launched in order to derive NASA polynomials that are suitable for use in a chemical mechanism. One species in particular whose thermodata is being updated in this way happens to be CO₂, which is part of the mechanism used in this study. As soon as the thermodata agent has deduced new thermodata for a species, the mechanism is updated. This then causes changes in the predicted atmospheric emission distributions.

Figure 6 shows concentration level contours overlaid on a Google map of The Hague, with some buildings. Distributions of carbon monoxide (CO, Figs. 6(a) and (b)) and unburnt hydrocarbons (uHC, Figs. 6(c) and (d)) are shown, each as predicted using the original mechanism by Wang et al. [44] and as predicted by the same mechanism with updated thermodata. It can be observed that the shape of the contours changes between the images for the original and the modified mechanisms. In addition, it should be noted in the legend that the range of concentration values corresponding to the colours of the contours also differs between the original and modified mechanisms. Taking this into account, the maximum predicted concentration level of CO decreases by about 3% from Fig. 6(a) to (b). Similarly, the maximum predicted concentration level of unburnt hydrocarbons changes by about 50% from Fig. 6(c) to 6(d).

5 Conclusions

We have developed agents that determine the reliability of the thermodata of sets of species based on error-cancelling balanced reactions, that conduct quantum calculations upon request via HTTP, and that deduce thermodynamic data from quantum calculations. These agents seamlessly fit into and extend the existing general-purpose knowledge-graph of the JPS. We have integrated these agents so that species whose thermodata have been identified as inconsistent automatically will be calculated at higher level of theory. Ther-

modata is automatically deduced from new quantum calculations, and propagated into an updated mechanism. As a proof of concept, and to demonstrate interoperability between heterogeneous data formats and software in a wider context, we have integrated this further into a cross-domain application that considers atmospheric pollutant dispersion simulations which utilise 3D geometries of buildings, and live weather data retrieved by agents from the world wide web. This use-case involves simulations across the length scales – from quantum calculations to macroscopic fluid flow. As the knowledge-graph grows, and the quality of the thermodynamic data within it improves, the mechanism used for the combustion simulation automatically updates. We have demonstrated in this multi-scale example how a new quantum calculation for a species can affect the distribution of pollutants in the atmosphere.

Acknowledgements

This work was partly funded by the National Research Foundation (NRF), Prime Minister’s Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. Markus Kraft gratefully acknowledges the support of the Alexander von Humboldt foundation. The authors are grateful to EPSRC (grant number: EP/R029369/1) and ARCHER for financial and computational support as a part of their funding to the UK Consortium on Turbulent Reacting Flows (www.ukctrf.com).

References

- [1] American Chemical Society. CAS Registry, 2019. URL <https://www.cas.org/support/documentation/chemical-substances>. Accessed 27 May 2020.
- [2] L. Atzori, A. Iera, and G. Morabito. The Internet of Things: A survey. *Computer Networks*, 54(15):2787–2805, 2010. doi:10.1016/j.comnet.2010.05.010.
- [3] T. Berners-Lee. Linked data – design issues, 2006. URL <https://www.w3.org/DesignIssues/LinkedData.html>. Accessed 10 June 2020.
- [4] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, May 2001. doi:10.1038/scientificamerican0501-34. URL <https://www.jstor.org/stable/10.2307/26059207>.
- [5] C. Bizer, T. Heath, and T. Berners-Lee. Linked Data – the story so far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009. doi:10.4018/jswis.2009081901.
- [6] P. Buerger, J. Akroyd, S. Mosbach, and M. Kraft. A systematic method to estimate and validate enthalpies of formation using error-cancelling balanced reactions. *Combustion and Flame*, 187:105–121, 2018. doi:10.1016/j.combustflame.2017.08.013.
- [7] Cambridge Environmental Research Consultants (CERC). Atmospheric Dispersion Modelling System (ADMS), 2020. URL <https://cerc.co.uk/environmental-software/ADMS-model.html>. Accessed on 17 March 2020.
- [8] CMCL Innovations. SRM Engine Suite, version 11.2.0, 2020. URL <https://cmclinnovations.com/solutions/products/srm/>. Accessed on 30 April 2020.
- [9] A. Devanand, I. A. Karimi, and M. Kraft. Optimal site selection for modular nuclear power plants. *Computers & Chemical Engineering*, 125:339–350, 2019. doi:10.1016/j.compchemeng.2019.03.024.
- [10] A. Devanand, G. Karmakar, N. Krdzavac, R. Rigo-Mariani, E. Y. S. Foo, I. A. Karimi, and M. Kraft. OntoPowSys: A power system ontology for cross domain interactions in an eco industrial park. *Energy and AI*, 1:100008, 2020. doi:10.1016/j.egyai.2020.100008.
- [11] A. Eibeck, M. Q. Lim, and M. Kraft. J-Park Simulator: An ontology-based platform for cross-domain scenarios in process industry. *Computers & Chemical Engineering*, 131:106586, 2019. doi:10.1016/j.compchemeng.2019.106586.
- [12] A. Eibeck, A. Chadzynski, M. Q. Lim, L. K. Aditya, L. Ong, A. Devanand, G. Karmakar, S. Mosbach, R. Lau, I. A. Karimi, E. Y. S. Foo, and M. Kraft. A parallel world framework for scenario analysis in knowledge graphs. *Data-Centric Engineering*, 2020. Accepted for publication.

- [13] F. Farazi, J. Akroyd, S. Mosbach, P. Buerger, D. Nurkowski, M. Salamanca, and M. Kraft. OntoKin: An ontology for chemical kinetic reaction mechanisms. *Journal of Chemical Information and Modeling*, 60(1):108–120, 2020. doi:10.1021/acs.jcim.9b00960.
- [14] F. Farazi, N. Krdzavac, J. Akroyd, S. Mosbach, A. Menon, D. Nurkowski, and M. Kraft. Linking reaction mechanisms and quantum chemistry: An ontological approach. *Computers & Chemical Engineering*, 137:106813, 2020. doi:10.1016/j.compchemeng.2020.106813.
- [15] F. Farazi, M. Salamanca, S. Mosbach, J. Akroyd, A. Eibeck, L. K. Aditya, A. Chadzynski, K. Pan, X. Zhou, S. Zhang, M. Q. Lim, and M. Kraft. A knowledge-graph approach to combustion chemistry and interoperability. *ACS Omega*, 2020. doi:10.1021/acsomega.0c02055. To appear.
- [16] D. Fensel, U. Şimşek, K. Angele, E. Huaman, E. Kärle, O. Panasiuk, I. Toma, J. Umbrich, and A. Wahler. *Knowledge Graphs – Methodology, Tools and Selected Use Cases*. Springer Nature, Switzerland, 2020.
- [17] M. Frenklach. Transforming data into knowledge – Process Informatics for combustion chemistry. *Proceedings of the Combustion Institute*, 31(1):125–140, 2007. doi:10.1016/j.proci.2006.08.121.
- [18] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, G. A. Petersson, H. Nakatsuji, X. Li, M. Caricato, A. V. Marenich, J. Bloino, B. G. Janesko, R. Gomperts, B. Mennucci, H. P. Hratchian, J. V. Ortiz, A. F. Izmaylov, J. L. Sonnenberg, D. Williams-Young, F. Ding, F. Lipparini, F. Egidi, J. Goings, B. Peng, A. Petrone, T. Henderson, D. Ranasinghe, V. G. Zakrzewski, J. Gao, N. Rega, G. Zheng, W. Liang, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, K. Throssell, J. A. Montgomery, Jr., J. E. Peralta, F. Ogliaro, M. J. Bearpark, J. J. Heyd, E. N. Brothers, K. N. Kudin, V. N. Staroverov, T. A. Keith, R. Kobayashi, J. Normand, K. Raghavachari, A. P. Rendell, J. C. Burant, S. S. Iyengar, J. Tomasi, M. Cossi, J. M. Millam, M. Klene, C. Adamo, R. Cammi, J. W. Ochterski, R. L. Martin, K. Morokuma, O. Farkas, J. B. Foresman, and D. J. Fox. Gaussian 16 Revision B.01, 2016. Gaussian Inc. Wallingford CT.
- [19] G. Fu, C. Batchelor, M. Dumontier, J. Hastings, E. Willighagen, and E. Bolton. PubChemRDF: Towards the semantic annotation of PubChem compound and substance databases. *Journal of Cheminformatics*, 7(1):34, 2015. doi:10.1186/s13321-015-0084-4.
- [20] C. W. Gao, J. W. Allen, W. H. Green, and R. H. West. Reaction Mechanism Generator: Automatic construction of chemical kinetic mechanisms. *Computer Physics Communications*, 203:212–225, 2016. doi:10.1016/j.cpc.2016.02.013.
- [21] C. F. Goldsmith, G. R. Magoon, and W. H. Green. Database of small molecule thermochemistry for combustion. *Journal of Physical Chemistry A*, 116(36):9033–9057, 2012. doi:10.1021/jp303819e.

- [22] J. Goodman. Computer software review: Reaxys. *Journal of Chemical Information and Modeling*, 49(12):2897–2898, 2009. doi:10.1021/ci900437n.
- [23] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical Reviews*, 136: B864–B871, 1964. doi:10.1103/PhysRev.136.B864.
- [24] R. D. Johnson III. NIST Computational Chemistry Comparison and Benchmark Database, NIST Standard Reference Database Number 101, Release 19, 2018. doi:10.18434/T47C7Z.
- [25] M. Keçeli, S. N. Elliott, Y.-P. Li, M. S. Johnson, C. Cavallotti, Y. Georgievskii, W. H. Green, M. Pelucchi, J. M. Wozniak, A. W. Jasper, and S. J. Klippenstein. Automated computational thermochemistry for butane oxidation: A prelude to predictive automated combustion kinetics. *Proceedings of the Combustion Institute*, 37(1):363–371, 2019. doi:10.1016/j.proci.2018.07.113.
- [26] S. Kim, P. Thiessen, E. Bolton, J. Chen, G. Fu, A. Gindulyte, L. Han, J. He, S. He, B. Shoemaker, J. Wang, B. Yu, J. Zhang, and S. Bryant. PubChem substance and compound databases. *Nucleic Acids Research*, 44(D1):D1202–D1213, 2016. doi:10.1093/nar/gkv951.
- [27] N. Krdzavac, S. Mosbach, D. Nurkowski, P. Buerger, J. Akroyd, J. Martin, A. Menon, and M. Kraft. An ontology and semantic web service for quantum chemistry calculations. *Journal of Chemical Information and Modeling*, 59(7):3154–3165, 2019. doi:10.1021/acs.jcim.9b00227.
- [28] V. R. Lambert and R. H. West. Identification, correction, and comparison of detailed kinetic models. In *9th US National Combustion Meeting*, 2015. URL <https://pdfs.semanticscholar.org/84bc/0933b0c29bdb7960e9106fcc51b6f024451e.pdf>.
- [29] H. Lasi, P. Fettke, H. G. Kemper, T. Feld, and M. Hoffmann. Industry 4.0. *Business & Information Systems Engineering*, 6(4):239–242, 2014. doi:10.1007/s12599-014-0334-4.
- [30] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, and C. Bizer. DBpedia – A large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web*, 6(2):167–195, 2015. doi:10.3233/SW-140134.
- [31] Y. P. Li, K. Han, C. A. Grambow, and W. H. Green. Self-evolving machine: A continuously improving model for molecular thermochemistry. *The Journal of Physical Chemistry A*, 123(10):2142–2152, 2019. doi:10.1021/acs.jpca.8b10789.
- [32] W. Marquardt, J. Morbach, A. Wiesner, and A. Yang. *OntoCAPE – A Re-Usable Ontology for Chemical Process Engineering*. Springer-Verlag Berlin Heidelberg, 1st edition, 2010. doi:10.1007/978-3-642-04655-1.
- [33] D. A. McQuarrie and J. D. Simon. *Molecular thermodynamics*. University Science Books, Sausalito CA, USA, 1999.

- [34] A. Menon, N. Krdzavac, and M. Kraft. From database to knowledge graph – using data in chemistry. *Current Opinion in Chemical Engineering*, 26:33–37, 2019. doi:10.1016/j.coche.2019.08.004.
- [35] S. Mosbach, M. S. Celnik, A. Raj, M. Kraft, H. R. Zhang, S. Kubo, and K.-O. Kim. Towards a detailed soot model for internal combustion engines. *Combustion and Flame*, 156(6):1156–1165, 2009. doi:10.1016/j.combustflame.2009.01.003.
- [36] N. F. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, and J. Taylor. Industry-scale knowledge graphs: lessons and challenges. *Communications of the ACM*, 62(8): 36–43, 2019. doi:10.1145/3331166.
- [37] M. Pan, J. Sikorski, C. A. Kastner, J. Akroyd, S. Mosbach, R. Lau, and M. Kraft. Applying Industry 4.0 to the Jurong Island eco-industrial park. *Energy Procedia*, 75:1536–1541, 2015. doi:10.1016/j.egypro.2015.07.313.
- [38] M. Pan, J. Sikorski, J. Akroyd, S. Mosbach, R. Lau, and M. Kraft. Design technologies for eco-industrial parks: From unit operations to processes, plants and industrial networks. *Applied Energy*, 175:305–323, 2016. doi:10.1016/j.apenergy.2016.05.019.
- [39] W. Phadungsukanan, M. Kraft, J. A. Townsend, and P. Murray-Rust. The semantics of Chemical Markup Language (CML) for computational chemistry : CompChem. *Journal of Cheminformatics*, 4(15):1–16, 2012. doi:10.1186/1758-2946-4-15.
- [40] R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1:140022, 2014. doi:10.1038/sdata.2014.22.
- [41] L. Ruddigkeit, R. van Deursen, L. C. Blum, and J. L. Reymond. Enumeration of 166 billion organic small molecules in the chemical universe database GDB-17. *Journal of Chemical Information and Modeling*, 52(11):2864–2875, 2012. doi:10.1021/ci300415d.
- [42] B. Ruscic, R. E. Pinzon, M. L. Morton, G. von Laszewski, S. J. Bittner, S. G. Nijssure, K. A. Amin, M. Minkoff, and A. F. Wagner. Introduction to Active Thermochemical Tables: Several “key” enthalpies of formation revisited. *The Journal of Physical Chemistry A*, 108(45):9979–9997, 2004. doi:10.1021/jp047912y.
- [43] K. R. Taylor, R. J. Gledhill, J. W. Essex, J. G. Frey, S. W. Harris, and D. C. De Roure. Bringing chemical data onto the Semantic Web. *Journal of Chemical Information and Modeling*, 46(3):939–952, 2006. doi:10.1021/ci050378m.
- [44] H. Wang, M. Yao, and R. D. Reitz. Development of a reduced primary reference fuel (PRF) mechanism for IC engine combustion simulations. *Energy & Fuels*, 27(12):7843–7853, 2013. doi:10.1021/ef401992e.
- [45] R. Weber and A. K. Wilson. Do composite methods achieve their target accuracy? *Computational and Theoretical Chemistry*, 1072:58–62, 2015. doi:10.1016/j.comptc.2015.08.015.

- [46] D. Weininger. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988. doi:[10.1021/ci00057a005](https://doi.org/10.1021/ci00057a005).
- [47] C. Zhang, A. Romagnoli, L. Zhou, and M. Kraft. Knowledge management of eco-industrial park for efficient energy utilization through ontology-based approach. *Applied Energy*, 204:1412–1421, 2017. doi:[10.1016/j.apenergy.2017.03.130](https://doi.org/10.1016/j.apenergy.2017.03.130).
- [48] L. Zhou, M. Pan, J. J. Sikorski, S. Garud, L. K. Aditya, M. J. Kleinlanghorst, I. A. Karimi, and M. Kraft. Towards an ontological infrastructure for chemical process simulation and optimization in the context of eco-industrial parks. *Applied Energy*, 204:1284–1298, 2017. doi:[10.1016/j.apenergy.2017.05.002](https://doi.org/10.1016/j.apenergy.2017.05.002).
- [49] L. Zhou, C. Zhang, I. A. Karimi, and M. Kraft. An ontology framework towards decentralized information management for eco-industrial parks. *Computers & Chemical Engineering*, 118:49–63, 2018. doi:[10.1016/j.compchemeng.2018.07.010](https://doi.org/10.1016/j.compchemeng.2018.07.010).
- [50] X. Zhou, A. Eibeck, M. Q. Lim, N. Krdzavac, and M. Kraft. An agent composition framework for the J-Park Simulator – a knowledge graph for the process industry. *Computers & Chemical Engineering*, 130:106577, 2019. doi:[10.1016/j.compchemeng.2019.106577](https://doi.org/10.1016/j.compchemeng.2019.106577).
- [51] X. Zhou, M. Q. Lim, and M. Kraft. A smart contract-based agent marketplace for the J-Park Simulator – a knowledge graph for the process industry. *Computers & Chemical Engineering*, 139:106896, 2020. doi:[10.1016/j.compchemeng.2020.106896](https://doi.org/10.1016/j.compchemeng.2020.106896).
- [52] T. Ziegler. Approximate density functional theory as a practical tool in molecular energetics and dynamics. *Chemical Reviews*, 91:651–667, 1991. doi:[10.1021/cr00005a001](https://doi.org/10.1021/cr00005a001).