A mixed quantum chemistry/machine learning approach for the fast and accurate prediction of biochemical redox potentials and its large-scale application to 315,000 redox reactions

Adrian Jinich[a*], Benjamin Sanchez-Lengeling[a*], Haniu Ren[a], Rebecca Harman[a], Alán Aspuru-Guzik[a,b]

[a]Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA, 02138
[b]Senior Fellow, Canadian Institute for Advanced Research, Toronto, Ontario M5G 1Z8, Canada
[*]These authors contributed equally to this work

Corresponding Author: Prof. Alán Aspuru-Guzik, Department of Chemistry and Chemical Biology, Harvard University, 12 Oxford Street, Cambridge, MA 02138. Phone: +1 (617) 384 8188. Email: alan@aspuru.com

Abstract

A quantitative understanding of the thermodynamics of biochemical reactions is essential for accurately modeling metabolism. The group contribution method (GCM) is one of the most widely used approaches to estimating standard Gibbs energies and redox potentials of reactions for which no experimental measurements are available. Previous work has shown that quantum chemical predictions of biochemical thermodynamics are a promising approach to overcome the limitations of GCM. However, the quantum chemistry approach is significantly more expensive. Here we use a combination of quantum chemistry and machine learning to obtain a fast and accurate method for predicting the thermodynamics of biochemical redox reactions. We focus on predicting the redox potentials of carbonyl functional group reductions to alcohols and amines, one of the most ubiquitous carbon redox transformations in biology. The method relies on semi-empirical calculations calibrated with Gaussian Process (GP) regression against available experimental data. Our approach results in higher predictive power than the GCM. We demonstrate the high-throughput applicability of the method by predicting the standard potentials of more than 315,000 redox reactions involving ~70,000 compounds, obtained with a network expansion algorithm that iteratively reduces and oxidizes a set of natural seed metabolites. We provide open access to all source code and data generated.

Introduction

All living systems are sustained by complex networks of biochemical reactions that extract energy from organic compounds and generate the building blocks that make up cells [1]. Recent work [2–5] has revived decades-old efforts [6–8] to obtain a quantitative understanding of the thermodynamics of such metabolic networks. Accurately predicting the thermodynamic parameters, such as Gibbs reaction energies and redox potentials, of biochemical reactions informs both metabolic engineering applications [9,10], and the discovery of evolutionary design principles of natural pathways [11]. This applies both to the prediction of the thermodynamics of known metabolic reactions and to non-natural reactions, which can be used to expand nature's metabolic toolkit [12,13].

Experimentally available reaction Gibbs energies and redox potentials provide coverage for only about ~10% of known natural metabolic reactions. The metabolic modeling community relies mainly on group contribution method (GCM) approaches [14–18] to estimate missing thermodynamic values. GCM decomposes metabolites into functional groups, and assign group energies by calibration against experimental data. The most widely used implementation of the GCM for biochemistry is the eQuilibrator [19,20], an online thermodynamics calculator. In addition to using GCM to predict the formation energies of compounds, the eQuilibrator makes use of experimental reactant formation energies when available, and combines them with group energies in a consistent manner in what is known as the component contribution method. This results in significant increases in accuracy [18]. However, whenever experimental reactant formation energies are not available, estimates are based solely on group contribution energies. Such GCM-based estimates, which do not capture intramolecular functional group interactions, have limited prediction accuracy [21].

In this work, we focus on predicting the thermodynamics of redox biochemistry. Redox reactions, which are fundamental to living systems and are ubiquitous throughout metabolism, consist of electron transfers between two or more redox pairs or half-reactions [22]. Quantum chemistry has recently emerged as an important alternative modelling tool for the accurate prediction of biochemical thermodynamics [21]. However, quantum chemical methods tend to have very high computational cost in comparison with the GCM or other cheminformatic-based alternatives.

Recent work in the intersection of quantum chemistry and machine learning has resulted in hybrid approaches that significantly lower computational cost without sacrificing prediction accuracy [23–26]. One such hybrid quantum chemistry/machine learning approach, previously applied to organic photovoltaics material design [27,28], relies on Gaussian process (GP) regression [29] to calibrate quantum chemical predictions against experimental data. GP regression is an established probabilistic framework in Machine Learning to build flexible models, that also furnishes uncertainty bounds on predicted data points. Gaussian process regression uses the

"kernel trick" [30] to make probabilistic predictions, leveraging the distance between a datapoint of interest and a training set.

Here we present a mixed quantum chemistry/machine learning approach for the accurate and high-throughput prediction of biochemical redox potentials. We focus on predicting the thermodynamics of carbonyl (C=O) functional group reductions to alcohols (C-O) or amines (C-N) (Figure 1A), two of the most abundant redox reaction categories in metabolism, and for which a significant amount of experimental data is available. The approach relies on predicting electronic energies of compounds using the semiempirical PM7 method (Figure 1B) [31], and then correcting for systematic errors in the predictions with GP regression against experimental data (Figure 1C, Figure 2). We innovate on existing GP-regression techniques by making use of novel reaction fingerprints [32] to compute the similarity between redox reactions. We demonstrate that the method results in higher predictive power than the GCM approach. Importantly, this comes at significantly lower computational cost than previously developed quantum chemistry approaches for biochemical thermochemistry. We demonstrate the high-throughput nature of the approach by predicting the redox potentials of what is to our knowledge the largest existing database of non-natural redox reactions, consisting of ~315,000 iterative oxidations and reductions of "seed" metabolites from the KEGG database (Figure 1A) [33–36]. We provide open access to the full source code and data, including metabolite geometries, electronic energies, pKa's and redox potentials.

Results

Calibration of semiempirical quantum chemistry redox potentials with GP regression yields high prediction accuracy at low computational cost.

Our goal is to develop a calibrated quantum chemistry modeling framework that can accurately predict the reduction potentials of biochemical redox reactions in a high throughput manner. Several different types of redox reactions that change the oxidation state of carbon atoms exist in biochemistry, including: reduction of carboxylic acids to aldehydes, reduction of carbonyls to alcohols or amines, and reduction of alcohols to hydrocarbons [4]. In this work, we focus our efforts on predicting the standard redox potentials (as a function of pH and ionic strength) for the reduction of carbonyls to alcohols or amines. This reaction category represents the most common type of carbon redox transformation in biochemistry [36]. In addition, it is the category for which the largest amount of experimental data is available in the NIST database for the thermodynamics of enzyme-catalyzed reactions (TECRDB) [37].

Our approach consists in running quantum chemistry calculations to obtain electronic structure energies of substrates and products for redox reactions interest. We take the difference in electronic energies as an estimate of the standard redox potential of the most abundant species (protonation states) of metabolites at pH=0. We use cheminformatic pKa estimates and the Alberty-Legendre transform [38] to convert the redox potentials to transformed standard redox

potentials $E^{o\prime}(pH = 7)$ (Methods). We then use GP regression to correct for systematic errors in the ab initio estimates (Figure 2) as well as pKa estimates that go into the model (Methods). In order to test and optimize the accuracy of our modeling framework, we compared predicted potentials against a dataset of 81 experimental redox potentials obtained from the NIST database of thermodynamics of enzyme-catalyzed reactions [37].

We tested several different model chemistries in order to maximize the accuracy and minimize the computational cost of our methodology (Methods and Supplementary Table 1). In this context, a model chemistry consists of the combination of multiple modelling choices that go into running a quantum chemical simulation. These include the conformer generation strategy and geometry optimization procedure, the electronic structure method (e.g. density functional theory (DFT), wave function approaches or semiempirical methods) as well as the water model used.

In addition, the Gaussian process regression-based calibration step depends on a particular choice of a covariance function, which in turn involves a distance or similarity measure between redox reactions or compounds. We found that the following combination of quantum model chemistry and kernel/distance function between reactions resulted in the most efficient prediction strategy: the PM7 [31] semiempirical method for both geometry optimizations and single point electronic energies, with COSMO implicit solvation [39], and a reaction fingerprint obtained by taking the difference between the Morgan fingerprint vectors of products and substrates [32], with a kernel function that is a mixture of a squared exponential kernel and a noise kernel (Methods).

Importantly, our GP-calibrated semiempirical quantum chemistry method results in higher predictive power than the most commonly used approach, the GCM (Figure 3). Using cross-validation to test prediction accuracy, GP-calibrated approach results in a higher pearson correlation coefficient (r), a higher coefficient of determination (R^2), and a lower mean absolute error (MAE) than GCM predictions when tested against the set of experimental potentials (Methods).

One of the advantages of our approach is its low computational cost (Supplementary Table 1) in comparison to other model chemistries. We note that the pipeline involves estimating pKas and major protonation states at pH=0 for all metabolites at a slight additional computational cost. However, this cost is also incurred by the GCM approach. Thus, the GP-PM7 approach is an accurate and computationally efficient modeling framework for the high-throughput prediction of biochemical redox potentials.

Generation of a large network of non-natural redox reactions obtained from iteratively reducing and oxidizing seed natural metabolites.

In order to demonstrate the high throughput nature of our method, we implemented a redox reaction network expansion algorithm to generate what is, to our knowledge, the largest database of non-natural biochemical redox reactions (Methods). The network expansion

algorithm makes use of the RDKit cheminformatics software to iteratively apply simplified molecular-input line-entry system (SMILES)[32,40] reaction strings to natural metabolites in the KEGG database. We generate two types of redox reactions, implemented in both the reductive and oxidative directions. The first type is the reduction carbonyls to alcohols (and the reverse oxidations); the second is the reduction of carbonyls to amines (and the reverse oxidations). Applying the cheminformatic transformation to metabolites in the KEGG database results in a first set of nearest-neighbor reduction or oxidation products. We then iteratively apply the reactions to the resulting set of products, until there are no more carbonyl functional groups to reduce (or alcohol or amine functional groups to oxidize).

Although the KEGG database of natural metabolic compounds and reactions contains molecules with up to 135 carbon atoms (a peptidoglycan cell wall component is the molecule with the highest number of carbon atoms [41]), we limit our network expansion algorithm to molecules with 21 carbon atoms or less. This threshold is imposed in consideration of the size of the compounds in the training set; if experimental data were available for larger compounds, this methodology could easily be adapted.

Table 1 shows the number and type of reactions that make up the network. It consists of ~75,000 compounds connected by ~315,000 reactions (reductions and oxidations). Of these reactions, 83% convert carbonyls to alcohols (or alcohols to carbonyls), while 17% convert carbonyls to amines (or amines to carbonyls). The large majority of the reactions (more than 80%) are oxidations of alcohols to carbonyls. Thus, the alcohol functional group is significantly more ubiquitous in the set of natural metabolites than either carbonyls or amine functional groups. The large majority of the reactions in the network - 97% - stem from the recursive *oxidations* of seed natural metabolites, while only ~3% come from recursive reductions. This is consistent with the notion that carbonyl functional groups can cause damage to macromolecules [42,43], and are thus kept at check in the cell.

We analyzed the structure of the resulting network by looking at the distribution of compound sizes. Figure 4a shows the number of compounds in the networks as a function of their number of carbon atoms. The distribution shows three peaks corresponding to molecules with 6n carbon atoms (n = 1,2,3). These peaks reflect the large number of products that result from combinatorially oxidizing all alcohol groups in mono-, di-, and trisaccharides (Figure 4B).

The distribution of predicted standard transformed redox potentials for the set of ~315,000 reactions generated with the iterative network expansion algorithm is shown in Figure 5. We use the millimolar standard state, E'^{m} instead of the more commonly used E'^{o} (where reactant concentrations are taken as 1M) since 1 mM is significantly closer to relevant physiological concentrations of metabolites in cells [4]. The resulting distributions for the two reaction categories peak at very close but slightly different values, with $\langle E'^{m} \rangle \sim -220$ mV for carbonyl to alcohol reductions and $\langle E'^{m} \rangle \sim -235$ mV for carbonyl to amine reductions.

Discussion

In this work we developed a mixed quantum chemistry/machine learning modeling approach for the prediction of biochemical redox potentials of carbonyl functional groups. The method is based on calculating the electronic energy difference between substrates and products in redox reactions using the semi-empirical quantum chemistry method PM7 [31]. The raw quantum chemistry estimates are then calibrated against experiment using GP regression.

We demonstrated that the method has better prediction power than the commonly used GCM approach. Furthermore, the computational resources required are significantly lower than previous quantum chemistry strategies for biochemical thermodynamics. The GCM decomposes compounds into discrete functional groups, and assigns energies to each group by calibrating via linear regression against experimentally-available Gibbs energies. It then estimates reaction energies as the group energy difference between products and substrates. As Figure 3A shows, GCM collapses a large fraction of the redox potentials of carbonyl compounds into a few set of values.

We note that in a recent variant of the GCM, group energies can be combined with experimental reactant Gibbs formation energies to significantly increase the accuracy of predictions, resulting in what's known as the CCM [18]. Since we seek to apply our modelling approach to predicting redox potentials of non-natural reactions, the CCM is not applicable to the vast majority of the reactions in our dataset.

Current work on GP calibration in computational chemistry relies on *molecular* fingerprints to compute the similarity between molecules of interest and compounds with experimental data. Our implementation is to our knowledge the first application of GP to reaction energy predictions using reaction fingerprints. We chose to use GP - as opposed to other machine learning techniques such as neural networks - based on the amount of experimental data available. GP are suited for this task, since they grows in complexity according to the data; a GP is robust to overfitting since it penalizes complex models, and is able to provide estimates of uncertainty on predictions [29,44,45]. In contrast, techniques such a deep neural networks - which have gained traction in recent machine learning applications to quantum chemistry [46–48] - are only applicable when much larger datasets are available.

We focused our approach to the prediction of the standard redox potentials of carbonyl functional groups (i.e. reduction to alcohols or amines). The application of GP regression to other redox transformations - including the reduction of carboxylic acids to aldehydes, and the reduction of alcohols to hydrocarbons - is limited by the availability of experimental data. This highlights the importance of a concerted effort to generate more experimental data for biochemical redox potentials in controlled conditions (such as pH, ionic strength, temperature and buffer) in order to increase the predictive power and scope of calibrated quantum chemical approaches.

We demonstrated the high-throughput nature of the methodology by generating a network of more than 315,000 non-natural biochemical reactions involving ~70,000 compounds. The set of carbonyl reductions to alcohols has an interesting long-tail distribution, suggesting that the molecular environment of the carbonyl can lead to significantly higher values of standard redox potential than an average natural carbonyl compound. Investigating the structure-energy relationships that explain this trend is an important next step. For instance, the carbonyl to alcohol reduction category displays a slight dependence of reduction potential on the number of carbonyl functional groups in the substrate compound (Figure S1).

Our algorithm is different from other network expansion algorithms previously applied to metabolic networks. By solely focusing on redox biochemistry, iterative application of the reactions converges to molecules that cannot be further reduced or oxidized. This contrasts with other network expansion algorithms [49,50] used to investigate origins-of-life scenarios, where several types of chemical transformations are considered, including carbon-bond formation and cleavage. Other network expansion algorithms also start from natural seed compounds but only consider natural metabolites [12]. In contrast, our algorithm uses natural metabolites as seeds but results in the expansion to a vast network of non-natural compounds.

We envision several applications of our biochemical redox potential prediction methodology. One such application is studying the thermodynamic landscape of specific families of natural metabolic compounds that undergo combinatorial reductions (or oxidations) of carbonyl (or alcohol or amine) functional groups. One such family of compounds are the brassinosteroids and the oxylipins [51–53], structurally diverse plant metabolites that play important roles in many physiological processes. In addition, in the context of drug metabolism, our methodology could be applied to obtain quantitative insights into the thermodynamics of redox transformations such as those mediated by the P450 (CYP) superfamily of enzymes [54].

To our knowledge, our database contains the largest set of geometric structures, electronic energies, pKa estimates, and redox potentials of natural metabolites and compounds related to these through oxidoreductive transformations. We make all the code and datasets generated in this work - including metabolite 3D geometries, electronic energies, pKa's and redox potential estimates - available as an open source repository.

Methods:

Generation of most abundant (major) protonation states at pH = 0

Our pipeline is based on performing electronic structure simulations of the most protonated species of each biochemical compound involved in every redox reaction of interest. We run the quantum chemistry simulations on the major protonation states at pH=0 to avoid errors associated to the prediction of anionic compound energies [21,55,56]. We use the Chemaxon calculator plugin (Marvin 17.7.0, 2017, ChemAxon) *cxcalc majormicrospecies* to generate, for

each substrate and product involved in a biochemical redox reaction of interest, the major microspecies at pH=0.

Geometry optimization

Using SMILES string representation [40] of the major species at pH=0 as input, we construct an initial 3D geometry using the Universal Force Field (UFF) as implemented in RDKit [57]. This geometry is then further refined by performing a geometry optimization using the PM7 [31] semiempirical method with the COSMO implicit solvation model [39] (see below for other geometry optimization model chemistries tested).

Single Point Energy estimates

We then compute the electronic energy $E_{electronic}$ of every compound involved in given a redox reaction using the PM7 semiempirical method [31] with the COSMO implicit solvation model [39] (see below for other single point energy SPE model chemistries tested). For both reductions of carbonyl groups to alcohols and amines, we add a hydrogen molecule to the substrate side of the redox reaction (in the direction of reduction) and compute its electronic energy. In addition, in order to balance reductions of carbonyls to amine functional groups, we add an ammonia molecule to the substrate side of the reaction and a water molecule to the product side of the reaction (in the direction of reduction), and compute their electronic energies.

We use the difference in the electronic energies of products and substrates, $\Delta E_{electronic}$, as an estimate of the standard redox potential E^o(major species at pH=0):

$$E^o(MS\ at\ pH = 0) \sim -\Delta E_{Electronic}/nF$$

where n is the number of electrons (2 for all reactions considered here) and F is Faraday's constant.

Our approach ignores the contribution of ro-vibrational enthalpies and entropies to Gibbs reaction energies. This significantly reduces the computational cost associated to quantum chemical simulations, and we correct for the systematic errors introduced by this approximation through the Gaussian process regression (see below). We note that previous work has shown that the value of $\Delta E_{electronic}$ obtained for these types of redox reactions strongly correlates linearly with the full ΔG^o_r prediction obtained from including ro-vibrational enthalpic and entropic contribution (data not shown).

Transforming chemical potentials E^o(major species at pH=0) to biochemical potentials, E'^o(pH = 7)

Having estimated the standard redox potential E^o(major species at pH=0) we use cheminformatic pKa estimates of reactants and the Alberty-Legendre transform [38] to convert E^o(major species at pH=0) to the *transformed* standard redox potential, E'^o(pH=7) which is a

function of pH [58]. In order to estimate the pKa's of all reactants in every redox reactions of interest, we use the chemaxon calculator plugin (Marvin 17.7.0, 2017, ChemAxon) *cxcalc pka*. Internally, the pKa calculator plugin is based on the calculation of partial charge of atoms in the molecule [59,60].

Gaussian process calibration using reaction fingerprints

We calibrate biochemical redox potentials E'°(pH = 7) obtained from quantum chemical simulations against available experimental data in order to correct for systematic errors in our simulations and the cheminformatic pKa estimates. One simple strategy of calibration of energy values would be to use a two-parameter linear regression. We note that our group has employed successfully this in the context of compounds for redox flow battery applications [61–66]. Here, we make use of the information provided by the difference between structural similarity between products and substrates and utilize a GP regression approach with reaction fingerprints to calibrate redox potential energies.

GP regression relies on a similarity or distance metric between data points. To construct the notion of similarity between reactions we utilize reaction fingerprints [32]. Several possible choices of kernel function and distance measure between molecules and reactions exist. Generally, the distance measures between molecules and reactions make use of fingerprint representations of compounds, which encode the structure of a molecule in a binary vector form [67]. Although several varieties of reaction fingerprints exist, ours are obtained from the difference of fingerprint vectors of products and substrates. For each molecule we generate a Morgan fingerprint [67], a fixed-length binary vector (bit vectors) indicating the absence (zeros) or presence (ones) of a particular graph-connectivity-environment. Each environment captures the local topological information of a molecule by mapping the local vicinity of connected atoms, along with their formal charges, type of chemical bond, and its position relative to a cyclic structure. When we consider the difference of two morgan fingerprints, we are looking at the subtraction and addition of molecular environments between two molecules (Figure S2). We expect that similar molecular transformations will have similar changes in molecular environments. When these vectors are normalized, the inner product between two rxn fingerprints x_1 and x_2 is the cosine distance [68], which is a measure of similarity between two vectors. This notion is adapted in Gaussian processes to quantify similarity between reactions in GP regression.

Predictions of GPs can be interpreted as weighted averages of the training data, where the weights are probabilistic in nature. GPs are distinct because of their associated covariance functions (i.e. the kernel) [30]. The equation below shows that the kernel we employ is a mixture of a squared exponential kernel and a noise kernel, which increases the robustness of the model.

$$K(x^*,\ x_i) = \sigma_{rxn} e^{-|x^* - x_i|^2/l} + \sigma_{noise} I$$

Here, K is the kernel function, σ_{rxn} is the variance of the reaction fingerprint kernel, l is the lengthscale parameter, and σ_{noise} is the white noise variance parameter.

Database of experimental redox potentials for calibration with Gaussian Process regression

To perform the calibration, use a dataset of available experimental transformed standard redox potentials. This dataset of redox half-reactions and their associated transformed standard reduction potentials $E^{o'}$(pH=7) was generated by Bar-Even et al. [4], by compiling experimental equilibrium constants (i.e. standard Gibbs reaction energies) from the NIST database of Thermodynamics of Enzyme-Catalyzed Reactions (TECRDB) [37] and the Robert Alberty database of biochemical compound Gibbs formation energies [38]. The dataset consists of 57 redox reactions that reduce a carbonyl functional group into an alcohol and 24 redox reactions that reduce a carbonyl functional group into an amine group.

To ensure we are not overfitting in the process of calibrating with GP regression, we used leave-one-out cross-validation (LOO-CV), where we train our model on all data points except one point and predict its value. This is done for all 81 data points, and reported predictions, accuracies, and resulting scatter plots always come from untrained data. It is also important to note that GPs are inherently robust to overfitting since the training procedure penalizes more complex models (higher-rank kernels) via an objective function [29].

Selecting a model chemistry with low computational cost and high prediction accuracy

In order to obtain fast and accurate predictions of redox potentials using the GP regression calibrated quantum chemistry strategy described above, we tested several different quantum model chemistries. A model chemistry consists of a combination of a geometry optimization (GO) procedure, a method to calculate the single point electronic energies (SPE) of optimized molecular geometries, as well as other modelling considerations, such as: number of geometric conformations per compound, water solvation model, and pKa estimation strategy.

We explored the following set of methods to obtain optimized geometric conformations: the Universal Force Field (UFF) [57], PM7 [31], DFTB [69], HF3-c [70] and PBEh-3c [71]. Additionally, to compute the electronic energies of the optimized structures through single point energy (SPE) calculations, we considered the following approaches: DFTB, PM7, HF3-c, PBEh-3c, DSD-PBEP86/SVP [72] and DLPNO-CCD(T)/SVP [73]. Given computational resource constraints, only a subset of all possible combinations of geometry optimization and SPE procedures were explored.

To select a model chemistry from the exploration set, we compared their prediction accuracies - after the GP regression calibration step described above - when tested on the set of experimental redox potentials. We quantified prediction accuracy using three different metrics: Pearson correlation coefficient (r), the coefficient of determination (R^2), and mean absolute error (MAE, in mV). Before applying the GP regression calibration step, the double-hybrid functional and linear-scaling couple-clustered methods DSD-PBEP86 and DLPNO-CCD(T) resulted in the

highest accuracy. Due to the relatively small variation in prediction accuracy, we picked the cheapest method that still gave reasonable accuracy (MAE < 30 meV), which turned out to be the PM7 semiempirical method for both geometry optimizations and single point electronic energies.

Group contribution method (GCM) estimates of redox potentials

We compared the accuracies obtained using the Gaussian progress regression calibrated quantum chemistry approach described above against those obtained with the commonly used GCM [14,18,74]. We use the GCM implementation of Noor et al. which was adapted to the thermodynamics of biochemical reactions [18]. Briefly, group energies of all compounds in the KEGG database are stored in a group matrix, with rows corresponding to compounds and columns corresponding to groups. Associated to each group is a group energy corresponding to the energy for the group's major protonation state (major species) at pH=7. In order to obtain a GCM-based estimate for a standard transformed reduction potential for the major species at pH = 7, we take the difference in group energy vectors for all products and substrate in the reaction. We then transform the resulting standard redox potentials E^o(major species at pH=7) to the standard *transformed* redox potential at pH = 7, $E^{o'}$(pH=7) using the same pKa estimates and Alberty transform approach described above.

Redox network expansion algorithm with cheminformatic reaction strings

We implemented a network expansion algorithm to iteratively reduce carbonyl functional groups and oxidized alcohol/amine functional groups for a subset of the natural metabolites of the KEGG database. Due to constraints in available computational resources to run electronic structure calculations, we used the set of all metabolic compounds in the KEGG database with 21 carbon atoms or less as seed metabolites. Using the RDKit open source cheminformatics software, we identify carbonyl functional groups (alcohol/amine functional groups) in each of the seed KEGG metabolites and transform them through a redox reaction to the corresponding reduced (oxidized) alcohol/amine (carbonyl). In practice, this is done using SMARTS reaction strings as implemented in RDKit, and iteratively applying them to SMILES representations [40] of the KEGG metabolites. The SMARTS reaction strings for the reduction and oxidation reactions considered here are shown below (Table 2 and Table 3)

We iteratively apply the reduction/oxidation reaction transformations to the compounds generated from KEGG seed metabolites. We terminate the iterations when the resulting product has, in the case of reductions, no more carbonyl functional groups that can be reduced (or in the case of oxidative transformations, no more alcohol/amine functional groups that can be oxidized to carbonyls). All software was written using the Python programming language.

Acknowledgements

We thank members of the Aspuru-Guzik lab for comments on the manuscript. The authors thank Harvard Research Computing for their support on using the Odyssey cluster. A.A.-G., A.J., and B.S.-L. Acknowledge support from SEAS^NVIDIA, Massively Parallel Programming and Computing (332986).

References

1. Nelson, D. L., Lehninger, A. L. & Cox, M. M. *Lehninger Principles of Biochemistry*. (Macmillan, 2008).
2. Noor, E. *et al.* Pathway thermodynamics highlights kinetic obstacles in central metabolism. *PLoS Comput. Biol.* **10,** e1003483 (2014).
3. Henry, C. S., Broadbelt, L. J. & Hatzimanikatis, V. Thermodynamics-based metabolic flux analysis. *Biophys. J.* **92,** 1792–1805 (2007).
4. Bar-Even, A., Flamholz, A., Noor, E. & Milo, R. Thermodynamic constraints shape the structure of carbon fixation pathways. *Biochim. Biophys. Acta* **1817,** 1646–1659 (2012).
5. Krumholz, E. W. & Libourel, I. G. L. Thermodynamic Constraints Improve Metabolic Networks. *Biophys. J.* **113,** 679–689 (2017).
6. Krebs, H. A., Kornberg, H. L. & Burton, K. A survey of the energy transformations in living matter. *Ergeb. Physiol.* **49,** 212–298 (1957).
7. Alberty, R. A. A short history of the thermodynamics of enzyme-catalyzed reactions. *J. Biol. Chem.* **279,** 27831–27836 (2004).
8. Burton, K. & Krebs, H. A. The free-energy changes associated with the individual steps of the tricarboxylic acid cycle, glycolysis and alcoholic fermentation and with the hydrolysis of the pyrophosphate groups of adenosinetriphosphate. *Biochem. J* **54,** 94–107 (1953).
9. Kiparissides, A. & Hatzimanikatis, V. Thermodynamics-based Metabolite Sensitivity Analysis in metabolic networks. *Metab. Eng.* **39,** 117–127 (2017).
10. Bar-Even, A., Noor, E., Lewis, N. E. & Milo, R. Design and analysis of synthetic carbon fixation pathways. *Proc. Natl. Acad. Sci. U. S. A.* **107,** 8889–8894 (2010).
11. Flamholz, A., Noor, E., Bar-Even, A., Liebermeister, W. & Milo, R. Glycolytic strategy as a tradeoff between energy yield and protein cost. *Proc. Natl. Acad. Sci. U. S. A.* **110,** 10039–10044 (2013).
12. Hadadi, N., Hafner, J., Shajkofci, A., Zisaki, A. & Hatzimanikatis, V. ATLAS of Biochemistry: A Repository of All Possible Biochemical Reactions for Synthetic Biology and Metabolic Engineering Studies. *ACS Synth. Biol.* **5,** 1155–1166 (2016).
13. Siegel, J. B. *et al.* Computational protein design enables a novel one-carbon assimilation pathway. *Proc. Natl. Acad. Sci. U. S. A.* **112,** 3704–3709 (2015).
14. Benson, S. W. & Buss, J. H. Additivity Rules for the Estimation of Molecular Properties. Thermodynamic Properties. *J. Chem. Phys.* **29,** 546–572 (1958).
15. Mavrovouniotis, M. L. Group contributions for estimating standard gibbs energies of formation of biochemical compounds in aqueous solution. *Biotechnol. Bioeng.* **36,**

1070–1082 (1990).
16. Mavrovouniotis, M. L. Estimation of standard Gibbs energy changes of biotransformations. *J. Biol. Chem.* **266,** 14440–14445 (1991).
17. Jankowski, M. D., Henry, C. S., Broadbelt, L. J. & Hatzimanikatis, V. Group contribution method for thermodynamic analysis of complex metabolic networks. *Biophys. J.* **95,** 1487–1499 (2008).
18. Noor, E., Haraldsdóttir, H. S., Milo, R. & Fleming, R. M. T. Consistent estimation of Gibbs energy using component contributions. *PLoS Comput. Biol.* **9,** e1003098 (2013).
19. Noor, E. *et al.* An integrated open framework for thermodynamics of reactions that combines accuracy and coverage. *Bioinformatics* **28,** 2037–2044 (2012).
20. Flamholz, A., Noor, E., Bar-Even, A. & Milo, R. eQuilibrator--the biochemical thermodynamics calculator. *Nucleic Acids Res.* **40,** D770–5 (2012).
21. Jinich, A. *et al.* Quantum chemical approach to estimating the thermodynamics of metabolic reactions. *Sci. Rep.* **4,** 7022 (2014).
22. Banerjee, R. *Redox Biochemistry*. (John Wiley & Sons, 2007).
23. Hansen, K. *et al.* Machine Learning Predictions of Molecular Properties: Accurate Many-Body Potentials and Nonlocality in Chemical Space. *J. Phys. Chem. Lett.* **6,** 2326–2331 (2015).
24. Ramakrishnan, R., Dral, P. O., Rupp, M. & von Lilienfeld, O. A. Big Data Meets Quantum Chemistry Approximations: The Δ-Machine Learning Approach. *J. Chem. Theory Comput.* **11,** 2087–2096 (2015).
25. Rupp, M., Tkatchenko, A., Müller, K.-R. & von Lilienfeld, O. A. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.* **108,** 058301 (2012).
26. Smith, J. S., Isayev, O. & Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **8,** 3192–3203 (2017).
27. Pyzer-Knapp, E. O., Simm, G. N. & Guzik, A. A. A Bayesian approach to calibrating high-throughput virtual screening results and application to organic photovoltaic materials. *Mater. Horiz.* **3,** 226–233 (2016).
28. Lopez, S. A., Sanchez-Lengeling, B., de Goes Soares, J. & Aspuru-Guzik, A. Design Principles and Top Non-Fullerene Acceptor Candidates for Organic Photovoltaics. *Joule* (2017). doi:10.1016/j.joule.2017.10.006
29. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning*. (MIT Press, 2006).
30. Kung, S. Y. *Kernel Methods and Machine Learning*. (Cambridge University Press, 2014).
31. Stewart, J. J. P. Optimization of parameters for semiempirical methods VI: more modifications to the NDDO approximations and re-optimization of parameters. *J. Mol. Model.* **19,** 1–32 (2013).
32. Schneider, N., Lowe, D. M., Sayle, R. A. & Landrum, G. A. Development of a novel fingerprint for chemical reactions and its application to large-scale reaction classification and similarity. *J. Chem. Inf. Model.* **55,** 39–53 (2015).
33. Kanehisa, M. *et al.* From genomics to chemical genomics: new developments in KEGG. *Nucleic Acids Res.* **34,** D354–7 (2006).
34. Kanehisa, M., Furumichi, M., Tanabe, M., Sato, Y. & Morishima, K. KEGG: new perspectives on genomes, pathways, diseases and drugs. *Nucleic Acids Res.* **45,** D353–D361

(2017).
35. Kanehisa, M., Sato, Y., Kawashima, M., Furumichi, M. & Tanabe, M. KEGG as a reference resource for gene and protein annotation. *Nucleic Acids Res.* **44,** D457–62 (2016).
36. Kanehisa, M. & Goto, S. KEGG: kyoto encyclopedia of genes and genomes. *Nucleic Acids Res.* **28,** 27–30 (2000).
37. Goldberg, R. N., Tewari, Y. B. & Bhat, T. N. Thermodynamics of enzyme-catalyzed reactions--a database for quantitative biochemistry. *Bioinformatics* **20,** 2874–2877 (2004).
38. Alberty, R. A. *Thermodynamics of Biochemical Reactions*. (John Wiley & Sons, 2005).
39. Klamt, A. & Schüürmann, G. COSMO: a new approach to dielectric screening in solvents with explicit expressions for the screening energy and its gradient. *J. Chem. Soc. Perkin Trans. 2* **0,** 799–805 (1993).
40. Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* **28,** 31–36 (1988).
41. van Heijenoort, J. Formation of the glycan chains in the synthesis of bacterial peptidoglycan. *Glycobiology* **11,** 25R–36R (2001).
42. O'Brien, P. J., Siraki, A. G. & Shangari, N. Aldehyde sources, metabolism, molecular toxicity mechanisms, and possible effects on human health. *Crit. Rev. Toxicol.* **35,** 609–662 (2005).
43. Ferguson, G. P. Protective mechanisms against toxic electrophiles in Escherichia coli. *Trends Microbiol.* **7,** 242–247 (1999).
44. Salimbeni, H. & Deisenroth, M. Doubly Stochastic Variational Inference for Deep Gaussian Processes. *arXiv [stat.ML]* (2017).
45. Gal, Y., van der Wilk, M. & Rasmussen, C. E. Distributed Variational Inference in Sparse Gaussian Process Regression and Latent Variable Models. *arXiv [stat.ML]* (2014).
46. Yao, K., Herr, J. E. & Parkhill, J. The many-body expansion combined with neural networks. *J. Chem. Phys.* **146,** 014106 (2017).
47. Wu, Z. *et al.* MoleculeNet: A Benchmark for Molecular Machine Learning. *arXiv [cs.LG]* (2017).
48. Faber, F. A. *et al.* Prediction Errors of Molecular Machine Learning Models Lower than Hybrid DFT Error. *J. Chem. Theory Comput.* **13,** 5255–5264 (2017).
49. Raymond, J. & Segrè, D. The effect of oxygen on biochemical networks and the evolution of complex life. *Science* **311,** 1764–1767 (2006).
50. Zubarev, D. Y., Rappoport, D. & Aspuru-Guzik, A. Uncertainty of prebiotic scenarios: the case of the non-enzymatic reverse tricarboxylic acid cycle. *Sci. Rep.* **5,** 8009 (2015).
51. Asami, T., Nakano, T. & Fujioka, S. Plant brassinosteroid hormones. *Vitam. Horm.* **72,** 479–504 (2005).
52. Fujioka, S. & Yokota, T. Biosynthesis and metabolism of brassinosteroids. *Annu. Rev. Plant Biol.* **54,** 137–164 (2003).
53. Mosblech, A., Feussner, I. & Heilmann, I. Oxylipins: structurally diverse metabolites from fatty acid oxidation. *Plant Physiol. Biochem.* **47,** 511–517 (2009).
54. Zanger, U. M. & Schwab, M. Cytochrome P450 enzymes in drug metabolism: regulation of gene expression, enzyme activities, and impact of genetic variation. *Pharmacol. Ther.* **138,** 103–141 (2013).
55. Tschumper, G. S. & Schaefer, H. F. Predicting electron affinities with density functional theory: Some positive results for negative ions. *J. Chem. Phys.* **107,** 2529–2541 (1997).

56. Simons, J. Molecular anions. *J. Phys. Chem. A* **112,** 6401–6511 (2008).
57. Rappe, A. K., Casewit, C. J., Colwell, K. S., Goddard, W. A. & Skiff, W. M. UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. *J. Am. Chem. Soc.* **114,** 10024–10035 (1992).
58. Alberty, R. A. *et al.* Recommendations for terminology and databases for biochemical thermodynamics. *Biophys. Chem.* **155,** 89–103 (2011).
59. Dixon, S. L. & Jurs, P. C. Estimation of pKa for organic oxyacids using calculated atomic charges. *J. Comput. Chem.* **14,** 1460–1467 (1993).
60. Csizmadia, F., Tsantili-Kakoulidou, A., Panderi, I. & Darvas, F. Prediction of distribution coefficient from structure. 1. Estimation method. *J. Pharm. Sci.* **86,** 865–871 (1997).
61. Er, S., Suh, C., Marshak, M. P. & Aspuru-Guzik, A. Computational design of molecules for an all-quinone redox flow battery. *Chem. Sci.* **6,** 885–893 (2015).
62. Huskinson, B. *et al.* A metal-free organic-inorganic aqueous flow battery. *Nature* **505,** 195–198 (2014).
63. Tong, L. *et al.* UV-Vis spectrophotometry of quinone flow battery electrolyte for in situ monitoring and improved electrochemical modeling of potential and quinhydrone formation. *Phys. Chem. Chem. Phys.* **19,** 31684–31691 (2017).
64. Yang, Z. *et al.* Alkaline Benzoquinone Aqueous Flow Battery for Large-Scale Storage of Electrical Energy. *Adv. Energy Mater.* doi:10.1002/aenm.201702056
65. Gerhardt, M. R. *et al.* Anthraquinone Derivatives in Aqueous Flow Batteries. *Adv. Energy Mater.* **7,** (2017).
66. Lin, K. *et al.* A redox-flow battery with an alloxazine-based organic electrolyte. *Nature Energy* **1,** nenergy2016102 (2016).
67. Rogers, D. & Hahn, M. Extended-connectivity fingerprints. *J. Chem. Inf. Model.* **50,** 742–754 (2010).
68. Bajusz, D., Rácz, A. & Héberger, K. Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? *J. Cheminform.* **7,** 20 (2015).
69. Elstner, M. *et al.* Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties. *Phys. Rev. B Condens. Matter* **58,** 7260–7268 (1998).
70. Sure, R. & Grimme, S. Corrected small basis set Hartree-Fock method for large systems. *J. Comput. Chem.* **34,** 1672–1685 (2013).
71. Grimme, S., Brandenburg, J. G., Bannwarth, C. & Hansen, A. Consistent structures and interactions by density functional theory with small atomic orbital basis sets. *J. Chem. Phys.* **143,** 054107 (2015).
72. Kozuch, S. & Martin, J. M. L. DSD-PBEP86: in search of the best double-hybrid DFT with spin-component scaled MP2 and dispersion corrections. *Phys. Chem. Chem. Phys.* **13,** 20104–20107 (2011).
73. Riplinger, C. & Neese, F. An efficient and near linear scaling pair natural orbital based local coupled cluster method. *J. Chem. Phys.* **138,** 034106 (2013).
74. Boudart, M. Thermochemical kinetics, 2nd. Ed., Sidney W. Benson, Wiley Interscience, 320 pp., $22.50, New York, 1976. *AIChE J.* **23,** 613–613 (1977).

Figure and Table Legends

Figure 1. Schematic representation of workflow. a) We generate of a network of redox reactions obtained from iterative reduction and oxidation of natural seed metabolites. The redox reactions considered are reductions of carbonyl functional groups to both alcohol and amine groups. b) Starting from a simplified molecular-input line-entry system (SMILES) [40] string representations of all the compounds in the network at pH=0, we pre-optimize molecular geometries using the Universal Force Field, and run semiempirical quantum chemistry calculations using PM7. c) We estimate the standard (millimolar state) redox potentials E^m for the major species at pH=0 from the difference in electronic energies. The transformed standard redox potentials $E^{m\prime}$(pH=7) is obtained from cheminformatic pKa estimates and the Alberty-Legendre transform. We calibrated a training set of compounds against available experimental $E^{m\prime}$(pH=7) values using Gaussian Process (GP) regression.

Figure 2. Schematic diagram of GP regression with reaction fingerprints. A redox reaction of interest x* is represented using a reaction fingerprint, which captures the molecular structure of both substrate (oxidized) and product (reduced). This reaction fingerprint is compared against the set of all redox reactions in the experimental dataset using the covariance function of the Gaussian process $k(x^*,x_i)$ (See Methods section). The GP correction ($E^{\circ\prime}_{GP}$) to the quantum chemical prediction ($E^{\circ\prime}_{PM7}$) is then obtained from a similarity-based average of the correction terms for the experimental dataset.

Figure 3. Comparison of prediction accuracy for group contribution method and quantum chemistry with GP regression. a) Scatter plot of experimental standard transformed redox potentials against GCM predictions. b) Scatter plot of experimental standard transformed redox potentials against GP-regressed quantum chemistry predictions. The experimental data consists of available standard transformed reduction potentials for carbonyl-to-alcohol and carbonyl-to-amine reductions. Data points in the scatter plot were obtained from a leave-one-out cross-validation (LOOCV) procedure, where the model is trained on all data points except the one point to be predicted. Note that we use the millimolar standard state, E'^m, where reactant concentrations are taken as 1 mM, since this is significantly closer to relevant physiological concentrations of metabolites in cells than 1 M.

Figure 4. The network of redox reactions obtained from iterative reductions and oxidations of natural metabolites. a) Carbon number distribution for the full set of ~70,000 compounds in the network. The peaks at multiples of 6-carbon atoms reflect the combinatorial oxidation of mono-, di-, and tri-saccharides and their multiple alcohol functional groups. b) Schematic diagram of a natural tri-saccharide being iteratively oxidized at all possible alcohol functional groups. Blue circles depict sites of an initial alcohol-to-carbonyl oxidations, and green circles depict a second oxidation site.

Figure 5. Distribution of predicted standard redox potentials (pH=7) for a set of ~315,000 non-natural redox reactions. Predicted potentials were obtained by running semi-empirical quantum chemical calculations (Methods) and correcting for systematic errors using Gaussian Process regression. The green distribution shows the E^{m}'(pH=7) values for 53,095 reductions of carbonyls to amines. The blue distribution shows the E^{m}'(pH=7) values for 261,903 reductions of carbonyls to alcohols.

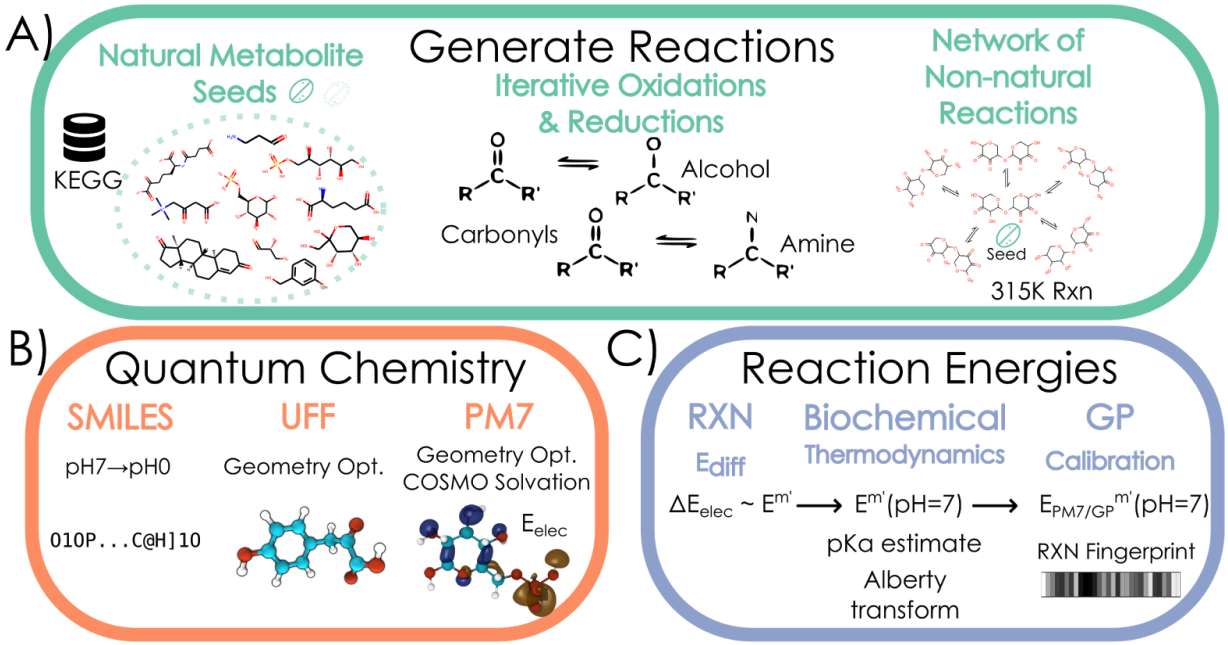
A)
Natural Metabolite Seeds
KEGG
Generate Reactions
Iterative Oxidations & Reductions
Alcohol
Carbonyls
Amine
Network of Non-natural Reactions
Seed
315K Rxn
B)
Quantum Chemistry
SMILES
UFF
PM7
pH7→pH0
Geometry Opt.
Geometry Opt. COSMO Solvation
O1OP...C@H]1O
E elec
C)
Reaction Energies
RXN
Biochemical
GP
Ediff
Thermodynamics
Calibration
ΔE elec ~ E m' → E m'(pH=7) → E PM7/GP m'(pH=7)
pKa estimate
Alberty transform
RXN Fingerprint

Figure 1

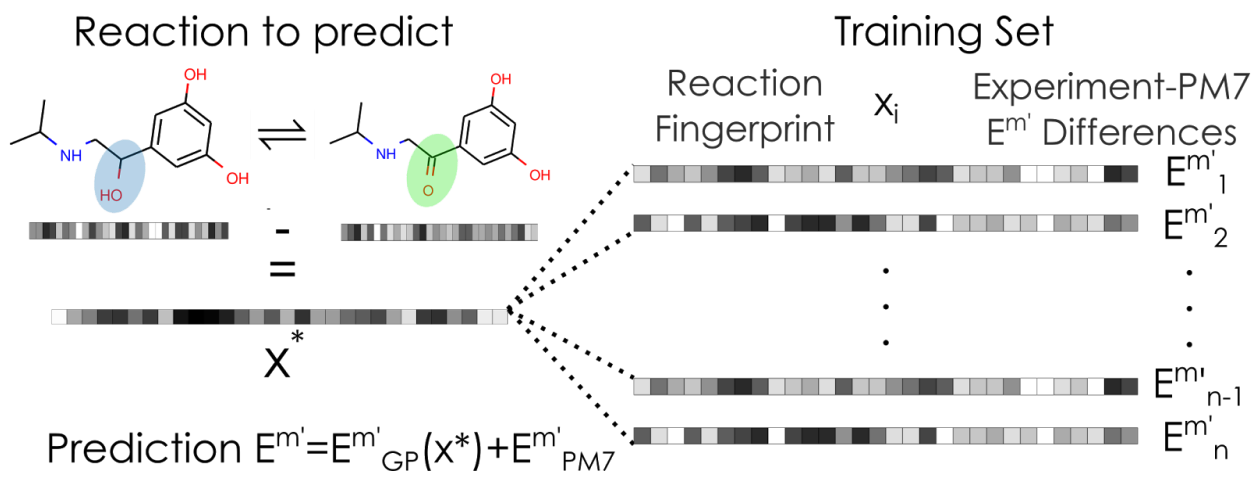
Reaction to predict
Training Set
Reaction Fingerprint
x_i
Experiment-PM7 $E^{m'}$ Differences
OH
NH
HO
OH
OH
NH
O
OH
-
=
x^*
$E^{m'}_1$
$E^{m'}_2$
$E^{m'}_{n-1}$
$E^{m'}_n$
Prediction $E^{m'}=E^{m'}_{GP}(x^*)+E^{m'}_{PM7}$

Figure 2

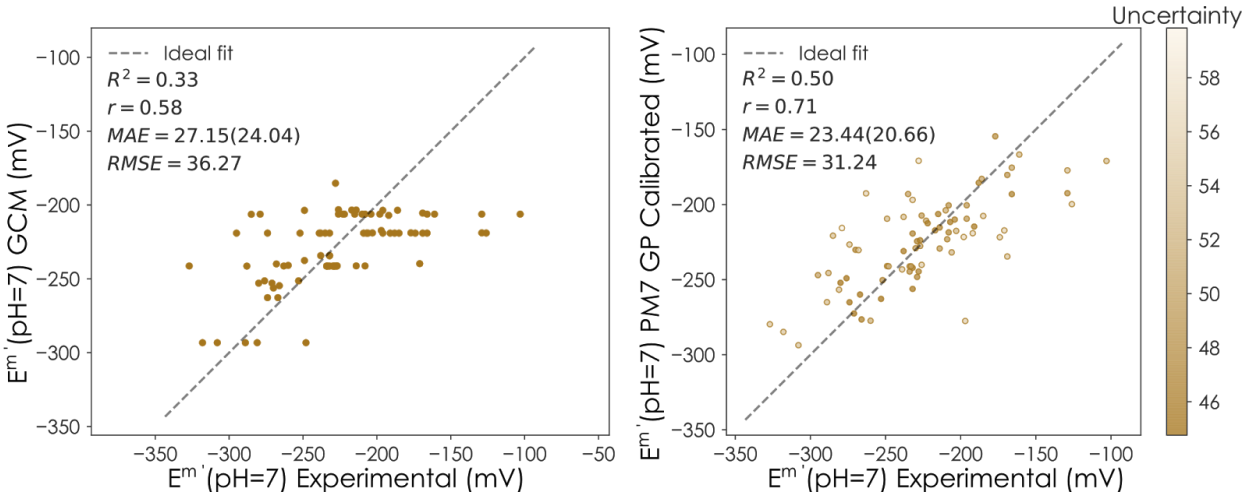
Ideal fit
$R^2 = 0.33$
$r = 0.58$
$MAE = 27.15(24.04)$
$RMSE = 36.27$
$E^{m'}$(pH=7) GCM (mV)
$E^{m'}$(pH=7) Experimental (mV)
Ideal fit
$R^2 = 0.50$
$r = 0.71$
$MAE = 23.44(20.66)$
$RMSE = 31.24$
$E^{m'}$(pH=7) PM7 GP Calibrated (mV)
$E^{m'}$(pH=7) Experimental (mV)
Uncertainty

Figure 3

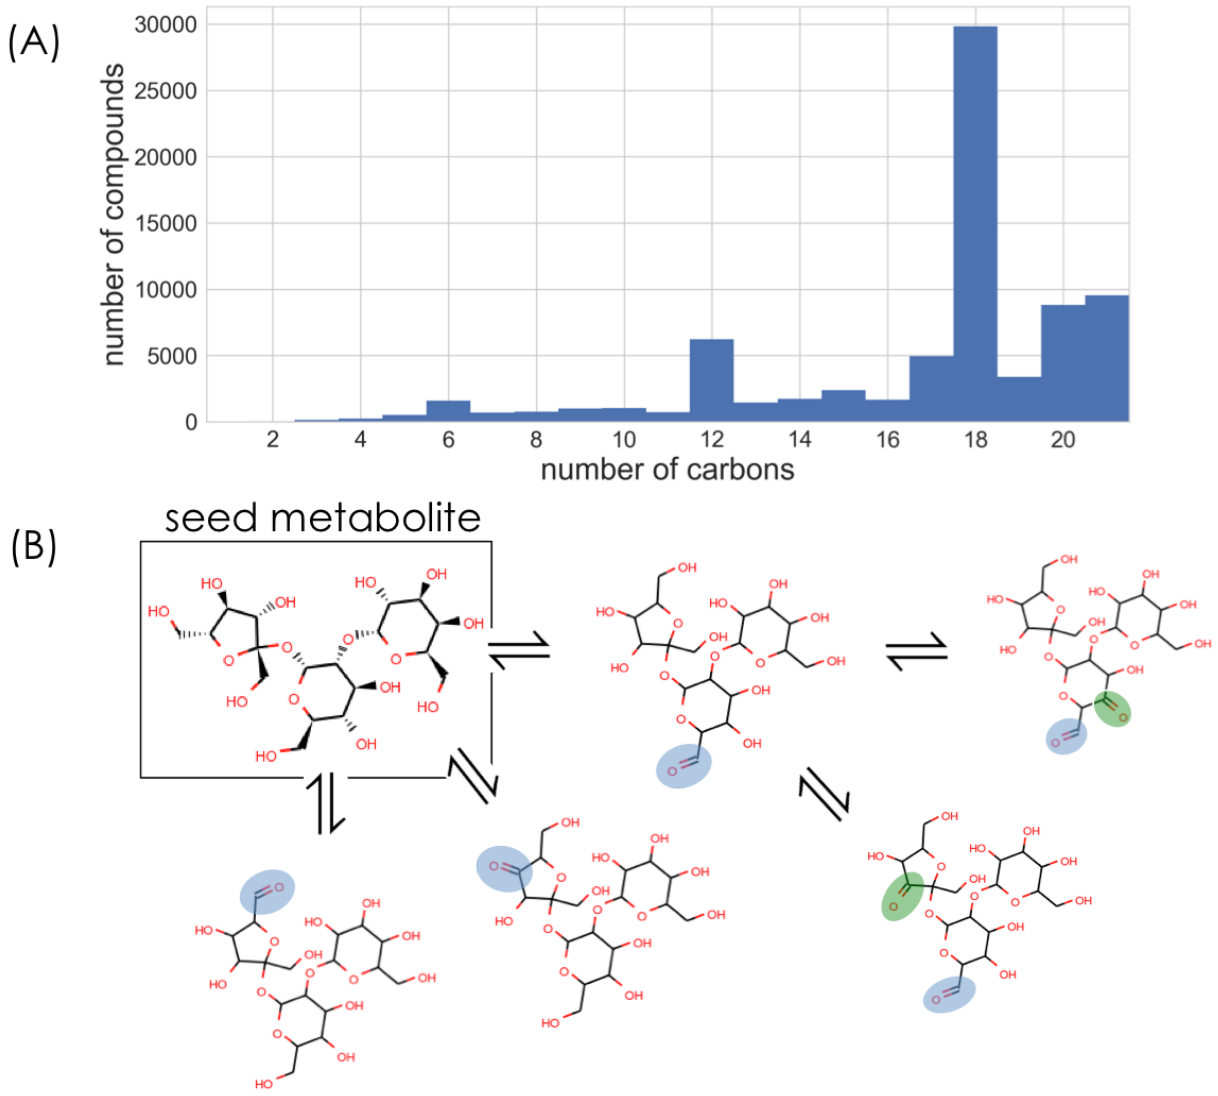
(A)
number of compounds
30000
25000
20000
15000
10000
5000
0
2
4
6
8
10
12
14
16
18
20
number of carbons
(B)
seed metabolite

Figure 4

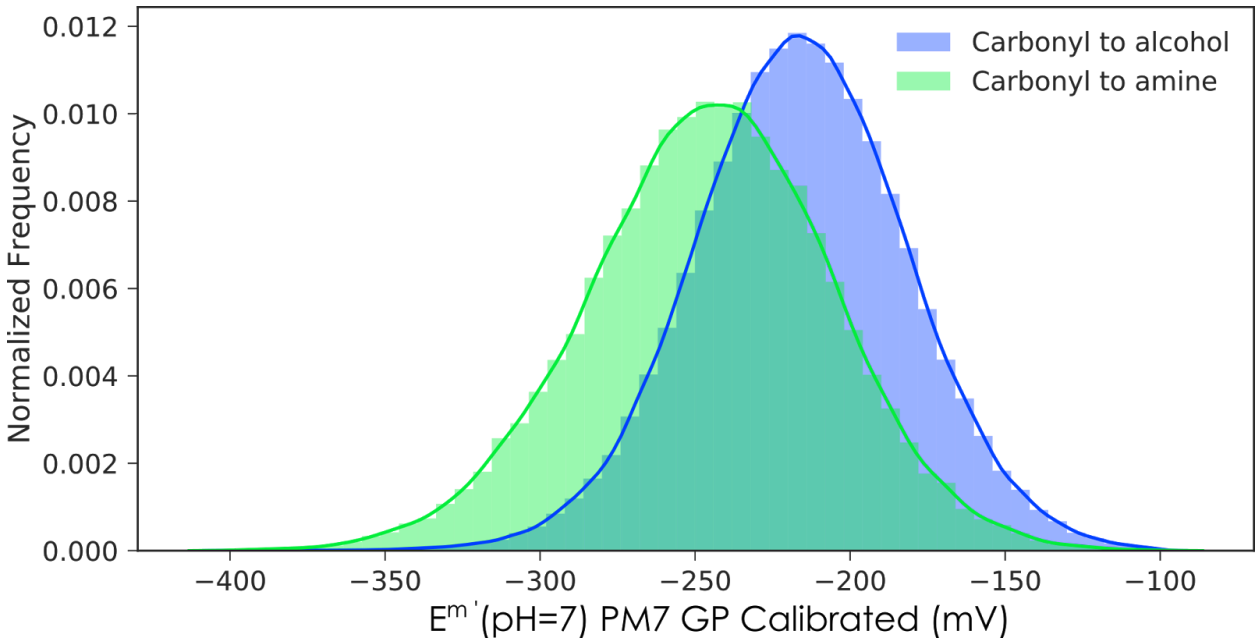
Carbonyl to alcohol
Carbonyl to amine
Normalized Frequency
0.012
0.010
0.008
0.006
0.004
0.002
0.000
−400
−350
−300
−250
−200
−150
−100
Em'(pH=7) PM7 GP Calibrated (mV)

Figure 5

Table 1: Number of reaction types in redox network

Reaction type	Number of reactions
carbonyl to alcohol reductions	4,299
carbonyl to amine reductions	4,299
alcohol to carbonyl oxidations	257,604
amine to carbonyl oxidations	48,796

Table 2: Cheminformatic reductions of carbonyls to alcohols or amines

Reaction Type	SMARTS reaction string
Aldehyde to alcohol	[CX3H1:2](=O)[#6:1] >>[#6:1][CX4H2:2][OX2H1]
Ketone to alcohol	[#6:1][CX3:2](=O)[#6:3] >>[#6:1][CX4H1:2]([#6:3])[OX2H1]
Aldehyde to amine	[#6:1][CX3H1:2](=O)>>[#6:1][CX4H2:2][NX3;H2;!$(NC=[!#6]);!$(NC#[!#6])]
Ketone to amine	[#6:1][CX3:2](=O)[#6:3]>>[C:1][CX4H1:2]([C:3])[NX3;H2;!$(NC=[!#6]);!$(NC#[!#6])]

Table 3: Cheminformatic oxidations of alcohols or amines to carbonyls

Reaction Type	SMARTS reaction string
Alcohol to aldehyde	[#6:1][CX4H2:2][OX2H1] >> [CX3H1:2](=O)[#6:1]
Alcohol to ketone	[#6:1][CX4H1:2]([#6:3])[OX2H1] >> [#6:1][CX3H0:2](=O)[#6:3]
Amine to aldehyde	[#6:1][CX4H2:2][NX3;H2;!$(NC=[!#6]);!$(NC#[!#6])]>>[#6:1][CX3H1:2](=O)
Amine to ketone	[C:1][CX4H1:2]([C:3])[NX3;H2;!$(NC=[!#6]);!$(NC#[!#6])]>>[#6:1][CX3H0:2](=O)[#6:3]