RE: Revisions to the manuscript “A compact model of Escherichia coli core and biosynthetic metabolism”

Dear Editor and Reviewers

We would like to thank you for taking the time to read and review our submission to the PCI Mathematical and Computational Biology. We are grateful for the detailed feedback provided on various aspects of the work, including the accompanying code, and have produced a revised version of the manuscript which, in our view, addresses the comments provided by the reviewers. Additionally, throughout the revision process, we have performed some minor stylistic changes, including a change in the color palettes used in the figures, which should improve the readability and interpretability of the work.

With the following letter, we provide answers to the reviewers’ comments and pinpoint how these have been addressed in the revised version of the article. In addition to the updated manuscript available in the ArXiv (https://arxiv.org/abs/2406.16596), we attach here an equivalent version wherein the changes introduced addressing reviewers’ comments have been marked in blue.

We hope that our updated manuscript answers all doubts and suggestions provided in the reviews, and look forward to hearing back.

Sincerely,

Marco Corrao,
on behalf of all the authors

Reviewer 1

The medium-scale metabolic model of E. coli that is created by the authors will likely be a useful resource for the metabolic modeling community. Although the work is not necessarily creative or novel, the goal of this research also doesn't require that. The work has been done carefully and is undoubtedly useful. The collection of various types of information about the metabolic model, as summarized in the 'functional annotation graph', will be useful as an information resource even if no metabolic modeling is performed. Finally, it is good to see that some commonly-used modeling methods are tried on the model, and that these turn out to be both feasible and that they give realistic results.

Since the work has been done carefully, I only have minor comments and suggestions on how to improve the current version, which I list below. In addition, I attach the filled-in survey as suggested by PCI to the reviewers.

We thank the reviewer for the supportive comments.

Minor comments and questions

In Section 2.2, the authors compare the behavior of the model by performing FBA with various substrate constraints, and by making "production envelopes". They show that the new model mostly agrees with the larger iML1515 model, and that where there is a deviation the new model is more reasonable. To me, this is slightly unconvincing since it is not clarified why the old model made the wrong predictions, and how this can be avoided in the future. Of course, it is not the authors' task to explain the mistakes in the larger model, but explaining these mistakes could give more guarantee that no similar mistakes are present in the iCH360 model.

We thank the reviewer for the comment. To better clarify, we have performed an extended investigation of the acetate-producing abilities of the two models both under aerobic and anaerobic conditions. As a result, Section 2.2 has been extended to include the results of this analysis and provide a rationale for why the metabolic routes used by iML1515 to produce more acetate are unlikely to be realistic under normal physiological conditions. Supporting our discussion, an additional Supplementary Figure (Supplementary Figure S7) was also introduced. Finally, we created a jupyter notebook detailing the analyses we performed and added it to the repository supporting the manuscript

(`Manuscript_Figures\notebooks\acetete_production_iCH360_vs_iML1515.ipynb`).

We would like to stress that these unrealistic predictions from iML1515 are not the result of "errors" in the metabolic model. Rather, they are the result of applying simple stoichiometric methods, such as FBA, to a large network with many degrees of freedom. It is possible that the inclusion of additional constraints, including thermodynamic feasibility or proteome allocation bounds, would automatically render such solutions unfeasible, but these are not always simple to implement. Our model iCH360 is smaller and thus filters out many of these behaviours *by construction*, providing users with a versatile and well-curated tool to investigate central and biosynthetic metabolism in *E. coli*. Clearly, this comes at the cost of limited applicability in other scenarios, such as those where the metabolic subsystems excluded by iCH360 (e.g. degradation pathways) are crucial to explaining or modelling a given phenotype. For such scenarios, the use of a genome-scale model (or an *ad hoc* reduction thereof), would still be the tool of choice.

As we believe this is an important point of clarification, we have added a paragraph in our discussion stressing this aspect.

In Section 2.4 the authors distinguish catalysis connections as primary and secondary. Connections are noted as secondary if the reaction "accounts only for negligible activity for the reaction in the wild-type strain". My question is if this will not be highly dependent on the conditions in which the wild-type strain is grown in the experiment. Can the authors provide some evidence that the experimental conditions that they take into account are broad enough, such that their primary/secondary annotations are applicable to almost all relevant conditions in which E. coli can grow?

The definition of primary and secondary relationships between reactions and enzymes is indeed of qualitative nature and heuristically defined. We agree with the reviewer in that, in a number of cases, the enzyme responsible for major activity for a reaction may be different depending on the growth condition, and that a more quantitative approach spanning many conditions where flux/expression data are available could improve this annotation and make it more universal. However, this is not a small endeavour and unfortunately out of scope for this manuscript. With this in mind, we have avoided introducing secondary catalysis relationships which were reported to be such only under very specific conditions and, as specified in our Methods, conservatively defaulted catalytic relationships to "primary" whenever sufficiently strong evidence was not available. Even then, our annotations are

based on evidence from the literature, which is often limited to one or few growth conditions. It is therefore not simple to provide an *a priori* justification for the validity of such annotations across all conditions. However, we believe that the validation we performed against an experimental dataset spanning a wide range of growth conditions is what ultimately justifies, in a statistical sense, the practical usefulness of this annotation. Had our annotations been highly condition-specific, we would not have observed significant clustering of fitness effects by the type of disrupted edges in the graph (even if, of course, false positives and negatives exist). In conclusion, we agree with the reviewer that this classification is not necessarily rigorous or comprehensive, but we believe that this information, with all its downsides, is still of practical value for model users. A good example, included in the manuscript, is the generation of unique reaction-enzyme mappings, often required for enzyme-constrained analyses.

"We did not find as significant differences" -> "We did not find any significant differences"

This was rephrased as suggested in the revised text.

The authors perform enzyme predictions using "Enzyme-constrained FBA", and state that their predicted enzyme abundances are quite good: "This analysis led to generally good predictions, with root mean squared error (RMSE, computed for log10-transformed enzyme abundances) ranging from 0.53 to 0.62 (Supplementary Figure 9)." Looking at Supplemental Figure 9, one can see that many enzymes are 10-fold (to even 50-fold) off. Although this statement is subjective, I think it is a bit of a stretch to call these "good predictions".

The phrase has been removed from the text.

Next, the authors fit the turnover numbers that they previously found to fit the enzyme data better, "To increase the predictive ability of the model". I don't think it is sufficiently made clear that the predictive ability of the model has been improved. Clearly, the fit gets better if hundreds of parameters can be adjusted. However, it is not clear whether this increases the predictive ability in other conditions. Maybe the authors could at least do a standard splitting of the data into a training and test dataset, and then seeing if fitting the data on the training set improves the predictive ability on the test set. The authors could do this by leaving one of the experimental conditions out.

We have now performed a leave-one-out cross-validation, as recommended by the reviewer, and extended Section 2.4 with the results of such analysis. Results (presented in the new Supplementary Figure 14) show that the predictions of enzyme abundances in a given condition are considerably improved even when data from that condition is not used for fitting. This supports the robustness of our parameter-adjusting procedure to condition-specific biases.

I don't understand the phrase "to be attributed to variability in enzyme demand across conditions", maybe it can be clarified.

With this sentence, we intended to specify that, since we fit a single apparent turnover number for all conditions considered, the residuals between measurements and our predictions (after parameter fitting) can be explained by variability in saturation levels for each enzyme across conditions, which a model with constant enzyme costs cannot capture. In the revised text, the expression "variability in enzyme demand" was changed to "variability in saturation levels", which we hope is clearer.

The authors mention that EFMs cannot be calculated for the full model due to redundancies. First of all, there are several alternative elementary pathway definitions that should make their enumeration feasible. I would, clearly, suggest calculating Elementary Conversion Modes (10.1016/j.patter.2020.100177) for this model, but a less biased overview of such alternative elementary pathways can be found in (https://doi.org/10.1371/journal.pcbi.1012472).

We thank the reviewer for the suggestion provided here. Indeed, we agree with them that alternative definitions of elementary pathways, including elementary conversion modes, is a valuable aspect to explore in relation to our model. Our intent here was not to make a general point on the virtues (or disadvantages) of EFM-based analyses. Rather, we intended to exemplify the use of our medium-scale model in the context of an analysis (EFM enumeration and screening) that would be complex, in the absence of specialised computing equipment, to perform at genome-scale.

In this sense, while the proposed analysis would certainly be valuable to perform, we believe that it is, at this moment, outside the scope of this manuscript. However, we have now

mentioned the potential of using elementary conversion modes as an alternative to EFMs in the Discussion, and noted the value of such an analysis as a future extension.

Second, when only redundant reaction pathways are deleted to produce the iCH360red model, can't knowledge of these deleted reactions be used to get the EFMs for the full model as well based on the EFMs that were calculated on the reduced model?

We believe that use of the word "redundancies" may have resulted in some misunderstanding. The reactions that were pruned from iCH360 to create iCH360red are redundant in the sense that other routes/reactions for the production of the same metabolite exist in the model (See supplementary table 4 for more details). However, these pruned reactions are *not* stoichiometrically identical to those maintained in the reduced model. Hence, to the best of our knowledge, enumerating the EFMs of iCH360 based on those of iCH360red is not necessarily different from the more general problem of enumerating the EFMs of a network given those of a subnetwork. To avoid confusion on this matter, we have replaced the word "redundancies" with "alternative metabolic routes" in the revised text.

When Muller et al. [28] are cited for proving that an optimal solution will only use one EFM, I think it is more correct to also cite Wortel et al. (10.1111/febs.12722) since these papers found this result simultaneously.

We now cite both papers.

"However, the satFBA formalism can also be used with additional flux bounds, thus going beyond EFM-based analysis." This is not true, it has been shown that the number of active EFMs is equal to the number of constraints (https://doi.org/10.1371/journal.pcbi.1006858). So by adding one or two additional constraints, the number of EFMs is still very limited, and one can easily find a small set of EFMs that give rise to the optimal flux distribution.

We appreciate and acknowledge the reviewer's comment that, even with additional flux bounds, the optimal solution could still be described by a small set of EFMs. Indeed, it was not our intention to present the satFBA analysis as a tool to overcome some limitations of EFM-based methods. Rather, we simply wanted to offer an alternative perspective on the same problem (modelling yield-growth trade-offs and their dependence on growth

conditions). In this light, the expression “thus going beyond EFM-based analysis” has been removed in the revised text.

I think adding a small analysis of this would be better, as the currently shown discrete increase in acetate production does not reflect the measured behavior from Basan et al. at all.

Indeed, the discrete increase in acetate production might seem as a limitation of the model when compared to the data from Basan et al. However, in that work, the method for manipulating the growth rate was changing the expression of the glucose transporter, while keeping the external concentration of glucose constant, where in our analysis the glucose concentration itself is gradually increased. Therefore, there is room for reallocation of resources between the uptake system and the rest of the proteome, something that indeed happens in our model. Furthermore, the model used in Basan et al. cannot be used to simulate a Monod curve as there is no representation of the external glucose concentration. Therefore, essentially these are distinct studies and we now believe that referencing the work of Basan et al. was a mistake. To avoid confusion, we removed these references in the revised manuscript.

The abbreviation “PTA” is introduced without being described.

We have fixed this typo with PMO (Probabilistic Metabolic Optimisation), which was introduced earlier in the text and also now listed in the *Abbreviations.*

Materials and Methods

Are the methods and analysis described in sufficient detail to allow replication by other researchers?

Not completely. The model is largely built by manual curation, and to me the choices made during that process are not entirely clear. I understand that it is hard to fully rationalize such manual curation steps, but some of the rationales could be more clearly presented with some examples.

For example, what type of reactions from the iML1515 model were left out, and why?

We have now extended Section 2.1 to state, besides the metabolic subsystems included in the model, the classes of reactions and pathways which we chose *not* to include in our model.

Reviewer 2:

The authors developed a central metabolic model of Escherichia coli, named iCH360, which expands on core carbon metabolism by including pathways for amino acid, fatty acid, and nucleotide biosynthesis. They demonstrated its utility through applications involving standard flux balance analysis (FBA) as well as integrated approaches with thermodynamic and kinetic calculations. With detailed annotations of enzymes and transporters, iCH360 serves as an excellent platform that integrates comprehensive biochemical knowledge about central metabolism. The model is exceptionally valuable from an educational perspective. One can envision using the reconstruction and applications described in this paper to develop a semester-long systems biology course, transitioning from simpler to more complex modeling approaches. From a research perspective, iCH360 is an ideal model for the initial testing of novel algorithms, such as those integrating gene expression or metabolite concentrations into metabolic networks. I anticipate that iCH360, along with its derivatives also provided by the authors, will become a standard tool in the metabolic network modeling community.

We thank the reviewer for the positive comments and outlook.

The manuscript is generally well-written, though some parts are overly complex. Below are some suggestions and questions to improve clarity:

1) In sections like 2.2, it would be beneficial to include comparisons with the ECC and ECC2 models alongside the parent model iML1515 to provide a clearer context for iCH360's advancements and distinctions.

In addition to the discussion presented in Section 2.1 related to the differences between iCH360 and ECC2, we now show production envelopes comparing iCH360, ECC and ECC2 in Supplementary Figure S9. As we discuss in the revised Section 2.2, the three models predict virtually identical maximal yields for the products considered, but some differences exist in their solution spaces closer to the maximal growth rates. However, as detailed in the revised text, the three models are equipped with different biomass requirements sourced from different parental models (hence predicting different biomass yields), which makes these differences expected.

2) Page 2: "For example, when designing and testing gene knockout strategies, genome-scale networks often wrongly predict unphysiological metabolic bypasses which have to be manually inspected and filtered out". This statement needs a supporting reference.

We have revised the text to include some concrete examples of unrealistic predictions produced by the genome-scale model using FBA, as well as citations to other articles where modifications in the genome-scale model had to be performed to filter out such unrealistic phenotypes. We stress, however, that these are only anecdotal examples and are not the results of an exhaustive investigation of the genome-scale parent predictions.

3) Figure 1: It is not clear what type of conversions are represented by the four reactions right above the grey area (the rectangular region with Phosphate-CO2-Ammonium-related pathways) on the lower left corner.

This information has been added to the figure.

4) Page 3: "The final assembled model (Figure 1) comprises 304 compartment-specific metabolites and 323 metabolic reactions mapped to 360 genes, thus qualifying as a medium-scale model ranging in between ECC and iML1515 (Supplementary Figure 5)." Please also indicate the number of unique metabolites excluding compartmentalization.

The number of chemically unique metabolites was included in the revised text.

5) Page 7: "By simultaneously fitting all available conditions, we ensured that our adjustment procedure is robust to condition-specific biases. Further, by introducing regularisation within our adjustment scheme, we penalised large deviations of parameters from the original dataset, increasing the robustness of the procedure to overfitting"

Based on these statements, can the authors do a "leave-one-condition-out" cross validation test to confirm the robustness of their procedure? I must add that I cannot see how this can be done based on the descriptions in Supplementary Information. If this indeed cannot be done, then the abovementioned statements may need to be revisited and clarified. For

example, what is the advantage of simultaneous fitting if each condition is fitted with a separate parameter?

We first would like to clarify that, for each enzyme-reaction pair, our procedure fits a single apparent turnover number, which we denote as the *typical* apparent turnover number, across all conditions considered. This condition-independent parameter is denoted as κ_i in Eq. 25 (Section A.5.1 of the Supplementary Information). In addition, for each condition considered, we also fit a scaling factor, denoted as τ_j in Eq. 25, which allows us to account for different total amounts of model enzymes measured in each condition. These scalings (a total of 8 parameters) merely amount to a vertical shift in the log-log plots in Figure 4 and would not be required if we were attempting to predict relative enzyme abundances, rather than absolute ones. We have stressed this aspect further in the revised text.

In light of these clarifications, it was in fact possible to run a leave-one-out cross-validation, and we did include this in our revised manuscript (see revised text in Section 2.4) and the new Supplementary Figure S14). Results of this analysis show that enzyme predictions for a condition are considerably improved even if data from that condition is not used for fitting, confirming the robustness of our procedure to condition-specific biases.

6) Page 8: The following sentence is not clear: "However, it is worth noting that this analysis was performed using a simple capacity-based enzyme cost function, which summarises different determinants of enzyme cost such as driving force and substrate saturation [27] in one single value, assumed to be constant. " Please rephrase.

Our intent was to specify that the cost function we used to compute the enzyme cost of each elementary flux mode assumes, for each reaction, a fixed enzyme investment per unit flux, regardless of the flux distribution considered. Hence, considering a more realistic enzyme cost function, as described in DOI:10.1371/journal.pcbi.1005167, could result in better estimates of the enzyme cost (and hence achievable growth rate) of each EFM, potentially improving this analysis. We have rephrased this sentence in the revised text, making this aspect more explicit.

7) Page 9: "As substrate availability is increased, the cost of substrate uptake decreases and higher growth rates are achieved by switching to lower yield, acetate secreting modes, in line with typical observations of overflow metabolism in E. coli [29, 30]. However, the satFBA formalism can also be used with additional flux bounds, thus going beyond EFM-based analysis. For example, if a positive lower bound on ATP hydrolysis is added as a maintenance requirement, optimal solutions to the satFBA problem will no longer be elementary modes, and the yield of the optimal solution no longer follows a piecewise constant profile (Supplementary Figure 16)"

This section is not clear. Specifically, the differences highlighted between Fig. 6 and Supp. Fig. 16, as mentioned in the text, are not readily apparent. Could the authors clarify this section and ensure the figures distinctly represent what is described?

With this paragraph (and related figure), we intended to highlight how introducing a constant energy requirement (in the form of a lower bound on the ATP-hydrolysis reaction ATPM), makes the optimal solution of the satFBA problem no longer an elementary mode. The two figures differ in that, without maintenance constraints (Figure 6, i.e. when the optimal solution is an EFM), the yield varies in a step-like manner as a function of external glucose concentration, where each “jump” corresponds to a change in optimal mode. When the maintenance constraint is introduced (Supplementary Figure S19), the yield profile no longer behaves this way. In the revised text, we have further clarified this aspect with additional text. Further, we have introduced zoomed insets in the yield plots, which should readily highlight the difference in yield profile between the two scenarios.

8) Figure 6C: "C: The yield of the optimal satFBA solution, computed as the ratio of biomass flux to glucose uptake, progressively decreases as external glucose availability increases."

Is the definition of yield the same everywhere in the manuscript? If so, this definition could be provided earlier to avoid confusion.

Indeed, our yields are always defined in this way and are expressed in [gDW/mol glucose], as described in the Methods. The phrase has now been removed from Figure 6C and moved earlier to Section 2.5, where the computation of yield is first introduced in the context of EFM analysis.

9) Supp.Tbl. 2: One type of edge between a reaction and a protein indicates inactive proteins. How many of the 360 genes in the model encode inactive proteins?

There are only two edges of this type and they point to the same protein (ilvG). The ilvG gene has a frameshift and encodes an inactivated protein. We have now included this information in the table. Note that the table is now updated to be Suppl Tbl 3.

10) Supp.Tbl. 2: The AND/OR example of the last row is not clear. How is this related to AND or OR connections?

In the knowledge graph we presented, logical nodes and edges can be used to define arbitrary logic in the model GPRs. Formally, a logical-OR node is used to create a logical expression where a number of arguments are ORed together, and the logical edges are what allows the specification of these arguments. Expanding on the example in the table, *E. coli* can use one of two thioredoxins, thioredoxin 1 (trxA) and thioredoxin 2 (trxC) (see e.g. 10.1093/emboj/17.19.5543), as a redox protein cofactor in a number of reactions (here, we will use TRDR as an example). We encode this relation in our graph by encoding both thioredoxins in a joint logical-OR node (THIOREDOXINS), which connects to the two thioredoxin proteins (RED-THIOREDOXIN-MONOMER and RED-THIOREDOXIN2-MONOMER) through logical edges. Hence, this node encodes the logical statement "RED-THIOREDOXIN-MONOMER OR RED-THIOREDOXIN2-MONOMER". Finally, the reaction is connected with the logical-OR node via a noncatalytic requirement relation, formally implying that either thioredoxin is required as a noncatalytic requirement of the reaction.

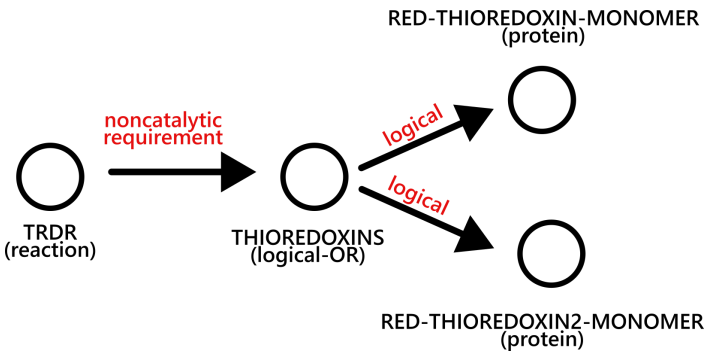

We have now added some text in the edge description table (which is indexed as Supplementary Table 3 in the revised manuscript) to clarify this aspect.

11)Supp.Inf. A.3: Eqn 18 and the bullet point that explains it are not clear. It would be useful to provide an example or examples from complicated GPRs with mixed ANDs and ORs.

We agree that the section was perhaps too condensed in the original manuscript. In this light, we have extensively revised this section, adding more context and focusing on more intuitive, biologically-grounded definitions rather than more formal logical language. Additionally, an illustrative schematic of how boolean rules are computed for different node types, depending on their children nodes, has been included as Supplementary Figure S22. We believe that, overall, these changes make the section much clearer to follow.

12)Page 15, Eqn. 5: Not clear where Eqn. 5 comes from. Also, mu is the specific growth rate, which, by definition, is the flux of biomass reaction (vBM). So it is not clear why we need the other terms and what is being achieved in this section (4.10).

In our EFM screening example, we had to quantify the specific growth rate achievable by each EFM. Since elementary flux modes represent directions in the flux space and can be scaled arbitrarily, they have to be normalised in some form in order to compute and compare the growth rate achievable by each of them. In our work, we performed such normalisation by scaling all EFMs to the same total enzyme cost (f_{enz} in Eq.5, whose value we estimated from data). In Eq.5, μ denotes the biomass flux upon this normalisation of the EFM (which is therefore independent of the scaling of the EFM), while v_{bm} denotes the biomass flux in the unscaled EFM (which is an arbitrary value that depends on the choice of normalisation used by the EFM enumeration algorithm). In other words, while μ and v_{bm} relate to the same physical quantity (a growth rate, in units of 1/h), the former is, by construction, a unique property of a flux *direction*, such as an EFM, while the latter is a property of a flux *vector* and is affected by its scaling.

13)Page 16: "where b is the min driving force (or the MDF after optimisation), c ∈ Rm is a vector of log-metabolite concentrations, lc; uc ∈ Rm are (log) lower and upper bounds on these concentrations, S ∈ Rm×N is the stoichiometric matrix of the model, R is the Boltzmann gas constant, and T is the temperature used for the computation of the free

energy estimates. " I believe what is said to be the Boltzmann gas constant here (traditionally this is denoted by a"k " not"R ") is actually the ideal gas constant.

We thank the reviewer for noting this mistake, which has been corrected in the revised text.

14) Section A.5.3: "Following the sMOMENT formulation of enzyme-constrained FBA [23], we consider a metabolic network with N reactions and M metabolites where all metabolic fluxes are positive (i.e. reversible reactions are split into forward and backwards components) and at most one enzyme is associated with each reaction." Based on the last part of this sentence, it is not clear what is done when multiple enzymes (isozymes) are involved in a reaction.

Our procedure for constructing unique reaction-enzyme mappings was previously outlined in the Methods (Construction of the enzyme-constrained metabolic model). To make this aspect clearer, a reference to that section has been included in this paragraph.

15) Section A.5.3, Eqn. 29: Biomass production reactions are typically irreversible. So, it is not clear what the negative biomass flux is needed to account for.

We believe that the confusion here relates to the fact that $\mathbf{r}_{eq}$ computed in Eq. 29, is the stoichiometry of the biomass reaction (i.e. a vector whose components are negative for the metabolites consumed by the reaction) and not the flux through it (which, indeed, is always positive). In the revised text, we have corrected the typeface of the $\mathbf{r}_{eq}$ (which is a vector and should have been written in bold) and included its dimensionality, making it clearer that it denotes a stoichiometry and not a flux.

16) Finally, could the authors please explain how much of the knowledge graph annotations are included in the JSON or SBML versions of the iCH360 model?

The boolean GPRs, which we computed using the knowledge graph, are included in the JSON/SBML versions of the iCH360 model. Additionally, the JSON/SBML versions of the enzyme-constrained model (EC-iCH360) contain additional annotation fields containing both the ID of the enzyme in the knowledge graph to which the reaction is mapped, its molecular weight, and the apparent turnover number (in 1/s) we used to compute the enzyme cost (see https://github.com/marco-corrao/iCH360/blob/main/Examples/enzyme_constrained_fba.ipyn

b for an example). The rest of the information contained in the graph, as detailed in the manuscript, is currently not part of the JSON and SBML models, and can be accessed and queried separately (see, for an example, https://github.com/marco-corrao/iCH360/blob/main/Examples/knowledge_graph.ipynb).

Reviewer 3:

The paper by Corrao et al introduces iCH360, a medium-scale metabolic model of Escherichia coli, and evaluates its predictive capabilities. The authors highlight that genome-scale models offer broad insights, but have limitations that smaller models can address. However, existing small-scale E. coli models often rely on algorithmic reductions and require extensive manual curation. To bridge this gap, this study aimed to create a well-curated medium-scale model.

iCH360 integrates diverse datasets, including gene-protein-reaction annotations, protein complex composition, thermodynamics, and kinetic constants. The model demonstrates predictive accuracy comparable to the genome-scale model, iML1515, while enabling elementary flux mode analysis, which is computationally infeasible for larger models. The study also applies thermodynamic and kinetic constraint methodologies, underscoring iCH360's potential for future applications.

The abstract presents a clear overview of the study. The introduction provides a good background on the developments made in modelling E. coli and provides a strong context for the introduction of the main motivations and aims for developing iCH360. The methods are comprehensive and provide a clear written description of the work completed. The assumptions made by the authors are clearly stated (in the methods, results or supplementary materials) and seem sound. Mathematical notation was supported with a well written commentary. Furthermore, scripts available online provide a comprehensive guide on the actual implementation of the methods. The discussion brings everything together and provides a broad summary of the results, whilst highlighting potential extensions and applications of iCH360. Throughout the paper, all references seemed relevant and accurate.

The authors invested great effort into writing a clear commentary and guide for executing the scripts. I downloaded the Zenodo repository (https://doi.org/10.5281/zenodo.13939696, on 2024-11-28) and installed the required packages. I had to install a number of additional packages and enclose the scripts I used for environment setup. They may be of use for future users of the repository. I was able to run all scripts (with some minor changes, see below) and reproduce all results and plots presented in the paper. I did not run the count_efs.m script due to the MATLAB requirement and instead used the efm_counts_unfiltered.csv file included in the repository.

We thank the reviewer for the positive outlook on our work. Particularly, we are grateful for the detailed feedback provided on our supporting code repository.

Some general points:

In my opinion, the findings could be reinforced in the abstract; that the predictions with iCH360 were in-line with iML1515 and/or experimental data; that the smaller scale of the model enabled the application of more advanced methodology; and that iCH360 is better than previous small-scale models. Additionally, considering the paper focuses on the development of iCH360, it should be named in the abstract.

We have revised the abstract to include the name of the model. However, we believe that quantitative comparisons between models with different origin and metabolic coverage come with a number of subtleties, and summarising them in a short sentence would be hard and potentially misleading. Hence, we preferred to leave analysis of these comparisons in the main text, where more in-depth discussion is possible.

The approach to determine an equivalent biomass reaction is attractive and well explained. There are two biomass reactions in iML1515: core and WT [1]. I assume that the WT biomass reaction was used but the authors should clarify this. Would the precursor fluxes change significantly when using the other reaction? Additionally, how were energetic requirements accounted for in iCH360 and were the chosen values inherited from iML1515? In addition, the stoichiometry of the equivalent biomass reaction should be made easily available as part of the supplementary material.

In fact, the “Core” biomass reaction in iML1515 was used to generate the equivalent biomass reaction in iCH360. This is the default objective used in the iML1515 model and its more parsimonious definition (relative to the WT counterpart) made it easier to manually

curate an extended model (which we denoted as iCH360ext in the manuscript) that was directly able to produce all its requirements, and hence generate the reference flux distribution required by our procedure. In practice, this last step can be performed using the genome-scale model itself to generate the reference flux distribution, but this wouldn't have allowed us to curate and control which pathways are used in the reference solution to produce each biomass requirement (which, in turn, affects the conversion into equivalent precursor cost). In light of this, while certainly valuable for future analysis, a direct and fair comparison between the equivalent costs computed based on the two biomass reactions would require the curation of a second extended model, which we believe to be outside of the scope of this manuscript. Nevertheless, we note that the "Core" biomass in iML1515 already converts certain complex biomass components present in WT biomass into equivalent costs of (relatively) simpler precursors. Thus, we can speculate that using the WT biomass wouldn't result in significant differences.

Addressing the rest of the comment, we have made the following changes in the revised manuscript (Section 4.1)

1. We have now specified which of the two biomass reactions was used to compute the equivalent costs,
2. We have explicitly stated that both growth and non-growth associated energy requirements were directly inherited from iML1515.
3. The stoichiometry of the equivalent biomass reaction has now been included as a Supplementary File (Supplementary File S1).

The authors highlight other small-scale models (ECC and ECC2) and include them in analysis of model properties. I believe further comparisons would be valuable in supporting iCH360. Is iCH360 better? How is it different? For example, simple predictions (FBA, phenotype phase planes) could be discussed. Whilst it is not a fair comparison for all carbon sources (only growth on glucose, glycerol, acetate and succinate were protected in ECC2 generation [2]), these comparisons could highlight the advantages of iCH360 being able to predict growth on a wider range of carbon sources.

The revised manuscript includes additional production envelopes (Supplementary Figure S9) comparing iCH360 with other small-scale models (ECC and ECC2), and a brief discussion of the observed similarities/differences in the main text (Section 2.2). As we discuss in the revised Section 2.2, the three models predict virtually identical maximal yields for the products considered, but some differences exist in their solution spaces closer to the maximal growth rates. As detailed in the revised text, the three models are equipped with

different biomass requirements sourced from different parental models (hence predicting different biomass yields), which makes these differences expected.

Relatedly, the authors could further investigate the consistency of predictions between iCH360 and iML1515. Whilst it does not necessarily validate the model predictions, the authors could compare fluxes (e.g. from pFBA) of reactions shared by both iCH360 and iML1515 across the different carbon sources investigated.

While this is a valuable (and relatively simple) analysis to perform, a thorough and fair discussion of the generated results would occupy substantial space in the manuscript, with several case-by-case investigations to be performed and discussed. Since, ultimately, we intend to present and showcase iCH360 as a standalone model, we believe that putting excessive emphasis on comparison with iML1515 would distract the reader from the main focus of the article.

In light of this, we have provided this analysis for readers that are interested in such comparison in a supplementary notebook of our supporting repository (`Analysis/iCH360_vs_iML1515/iCH360_vs_iML1515_pFBA_fluxes_comparison.ipynb`). We have reminded the reader of the additional analyses available in the supporting repository in the Code Availability section of the manuscript.

Where possible, the sources/literature used for defining the different types of catalytic edges should be cited or made available in a supplementary file.

This information, which was previously present in the supporting repository and as metadata in the knowledge graph, is now also included as Supplementary File S2.

The authors approximate enzyme abundance by constructing an augmented matrix $\hat{E}$, allowing one to account for polypeptides that are part of additional enzymatic complexes which are not part of the model. How often are these polypeptides mapped out of the model? Based on the polypeptides and the in-model complexes they map to, are there certain types of complexes particularly affected by this and is this of interest?

There are 7 polypeptides in the model that also map to out-of-model enzyme complexes. We have included this information in the revised text. These are a rather small share of the total number of polypeptides in the model and we could not identify any striking trend in the type of complexes they belong to.

Finally, the authors could expand on their interpretation of the results from the thermodynamic-based analysis.

Indeed, we have described the results of the thermodynamic analysis only briefly. The main purpose of running it was to verify that the thermodynamic parameters (i.e. the vector of reaction standard Gibbs energies described by a multivariate distribution) can technically be used (e.g. in thermodynamic-based flux balance analysis) by checking that the MDF score for the pFBA solutions is positive, and that PMO predictions are within reasonable ranges. The last part of the analysis (comparing between the predicted flux and the flux-force efficacy) might indeed shed some light on design principles in metabolism, and we have added the following sentence:

"Establishing a causative link can be very difficult, but we can speculate that this negative correlation might be explained by metabolism evolving to compensate for the low efficacy of some reactions with a higher expression level of their catalysing enzyme (Noor et al. 2014,)."

We feel that this issue has been sufficiently dealt with in other publications.

Overall, this is a solid piece of work, that is well-written and with good accompanying code. With the additional validations mentioned, I would agree with the authors that "iCH360 holds the potential to become a reference metabolic model for E. coli ".

References

[1] Jonathan M. Monk et al. "iML1515, a Knowledge base That Computes Escherichia Coli Traits". In: Nature Biotechnology 35.10 (Oct. 2017), pp. 904–908. DOI: 10.1038/nbt.3956. URL: https://www.nature.com/articles/nbt.3956.

[2] Oliver Hädicke and Steffen Klamt ."EColiCore2: A Reference Network Model of the Central Metabolism of Escherichia Coli and Relationships to Its Genome-Scale Parent Model". In: Scientific Reports 7.1 (Jan. 3, 2017), p. 39647. DOI: 10.1038/srep39647. URL:https://www.nature.com/articles/srep39647.

A number of minor issues found when running the scripts:

1. primary_seondary_counting.Rmd was empty. I assume it is an unused file.

2. In enzyme_allocation_predictions.ipynb, the path is written as /manuscript_figures/ but should be /Manuscript_Figures/.

3. The path to the directory Knowledge_Graph was written as Knowledge_graph in generate_model_tables.ipynb, catalytic_disruption_analysis.ipynb and estimate_enzyme_abundances_from_pp_counts.ipynb.

4. Flux-force-efficacy_vs_measured_enzyme_abundance.R and MDF_PTA.ipynb initially failed due to usage of an unavailable file ../../Analysis/PTA/out/pta_reactions_data.csv. Replacing the path with ../../Analysis/PTA/out/pta_fluxes.csv fixed this issue and allowed for results and figures to be produced.

5. For compute_efm_cost_yield.ipynb to work, I had to create a directory mkdir Analysis/EFM_growth_yield_screening/out.

We thank the reviewer for thoroughly inspecting our supporting repository and spotting these mistakes in the paths, which have likely occurred upon refactoring the repository prior to submission. These have been fixed in the most recent versions of the supporting repository on Zenodo and Github.

Environment setup

conda create -n ich360

conda activate ich360

conda config --add channels conda-forge

conda config --set channel_priority strict

export GRB_LICENSE_FILE=/path/to/gurobi/11.0.1/gurobi.lic

export GUROBI_HOME=/path/to/gurobi/11.0.1/

export PATH="${PATH}:${GUROBI_HOME}/bin"

```
export LD_LIBRARY_PATH="${LD_LIBRARY_PATH}:${GUROBI_HOME}/lib"
conda install -c conda-forge -c gurobi python==3.9.20 cobra==0.29.0 \
numpy==1.24.1 scipy==1.10.1 pandas==1.5.3 matplotlib==3.7.1 \
seaborn==0.12.2 networkx==3.0 tqdm==4.65.0 requests==2.28.2 \
casadi==3.6.3 cvxpy==1.5.2 equilibrator-api==0.4.7 nb_conda_kernels==2.5.1 \
notebook==7.1.3 pyvis==0.3.1 r-base==4.3.3 r-dplyr==1.1.4 r-ggplot2==3.5.1 \
r-reshape2==1.4.4 r-ggsci==3.2.0 r-ggpubr==0.6.0 rpy2==3.5.11 gurobi==11.0.1 \
adjusttext==1.3.0 r-svglite==2.1.3
pip install equilibrator-assets==0.4.1 efm tool==0.2.1 enkie==0.1.3 \
straindesign==1.13 pta==0.6.0
```

In addition to setting up the conda environment, there were a number of additional steps needed. I include them here for reference.

1. For first time use of enki and to run pta.ipynb, I needed to create a directory for it in my cache mkdir /lisc/user/coltman/.cache/enkei. This issue was previously reported https://gitlab.com/csb.ethz/enkie/-/issues/1

2. Due to the versions fixed by the package requirements, there were some issues with the setup of equilibrator. I had to manually save the files from the following repos (https://zenodo.org/records/4128543, https://zenodo.org/records/4013789, https://zenodo.org/records/4010930) to my ~/.cache/equilibrator in order to run drg0_estimation.ipynb

We thank the reviewer for flagging these issues in the installation of required packages. We have now added a note in the main *readme* of the repository pointing out that these steps may be necessary to run the scripts correctly.