Towards Using Large Language Models and Deep Reinforcement Learning for Inertial Fusion Energy

Vadim Elisseev
IBM Research
University of Oxford
Wrexham University
vadim.v.elisseev@ibm.com

Massimiliano Esposito
IBM Research
Massimiliano.Esposito@ibm.com

James Sexton
IBM Research
sextonjc@us.ibm.com

Abstract

Fusion energy research has long captured the public imagination for its applications to fundamental physics, material sciences, and as a low-carbon-footprint electrical power source. The National Ignition Facility (NIF) recently demonstrated that focusing lasers onto a very small target of hydrogen isotopes can produce conditions for nuclear fusion. Despite such remarkable progress, sustainable production of inertial fusion energy (IFE) still presents a huge challenge due to a vast space of parameters that must be explored in order to find optimum conditions for a thermonuclear ignition. It is perceived that artificial intelligence (AI) can play a crucial role in advancing IFE technology. We present our vision of how large language models (LLM) and deep reinforcement learning (DRL) can guide IFE research.

1 Introduction

In a fusion reaction, nuclei of the two isotopes of hydrogen, deuterium (D) and tritium (T), are forced together by high temperature and pressure and fuse to form a helium nucleus, also known as an α particle releasing some of the mass of the hydrogen as energy Chen [2016]. Inertial confinement fusion (ICF) Nuckolls et al. [1972] uses lasers to implode a capsule with DT fuel. Multiple laser beams (192 in case of NIF Facility [2024]) are focused on a millimetre-size target and deliver several mega joules (MJ) of energy transforming target surface into plasma, which accelerates outwards and creates an implosion accelerating the target shell inwards compressing the inner core and increasing its temperature to several kiloelectronvolts (KeV) producing DT fusion reactions. If density of the compressed core is high enough, α particles deposit their energy within the hot core, amplifying the fusion reaction rates. Under the appropriate conditions of pressure, temperature and confinement time, this process becomes intense enough to ignite the central hot spot, producing a thermonuclear burn wave propagating through dense fuel, yielding fusion energy many times higher than lasers energy input Gopalaswamy et al. [2019].

Experiments by the National Ignition Facility (NIF) have recently achieved ignition Abu-Shawareb et al. [2022], Facility [2024] and provided fusion yield measured as *fusion energy generated / input energy* of up to 2.5. Much higher yields are required for sustainable IFE production. Ability to achieve those depends on many factors related to parameters of the targets (fuel mix, material and geometry, manufacturing imperfections, etc.), spatial and temporal characteristics of laser beams as well as on various parametric instabilities, which can occur during laser-plasma interactions and

Machine Learning and the Physical Sciences Workshop, NeurIPS 2024.

prevent achieving of ignition conditions Tikhonchuk [2013]. Further advancements in IFE research therefore imply ability to explore a huge space of parameters that must be navigated in order to find their optimum configuration.

We present our vision of using large language models (LLMs) and deep reinforcement learning (DRL) using to accelerate research of such complex physical phenomena as IFE.

The rest of the paper is organized as follows: Section 2 describes relevant work in the area of fusion, including magnetic confinement fusion (MCF) DOE [2023], which is another way to produce fusion energy. Section 3 outlines our vision, Section 4 describes challenges that we are aiming to address, Section 5 talks about implementation of our vision, and Section 6 provides concluding remarks.

2 Related work

An example of using reinforcement learning (RL) for plasma control in fusion was presented in Jonas Degrave [2022], where the authors employed Maximum a Posterior Policy Optimization (MPO), an actor-critic algorithm, to manage plasma in a Tokamak reactor Hofmann et al. [1994]. In this study, domain experts defined the objectives that plasma must achieve, which included various intermediate targets and terminal states. These objectives were encapsulated into a carefully designed reward function, which, in conjunction with a predefined set of actions and a simulation environment, was used to guide the RL algorithm in learning optimal control strategies. The selection of MPO was driven by the need for sample efficiency, as the computational cost of the simulation environment was prohibitively high.

Another significant contribution in this field can be found in Seo et al. [2024], where the authors utilized Deep Deterministic Policy Gradient (DDPG) to optimize plasma pressure up to the threshold of instability. The RL model operated using coarse control actions, such as total beam power and plasma triangularity, which were subsequently refined by a plasma control system that managed the magnetic coils and the power distribution of individual beams to meet the high-level control objectives set by the AI.

While these two works have shown incredible results they come with limitations like the need of using tailored experts rewards to achieve state of the art results, the necessity of adding human devised sub-tactical goals, in the case of adopting large state spaces, which imply costly simulations and a need for using techniques like MPO, which might not be always the most performing ones.

Very recently, scholars have begun investigating the use of large language models (LLMs) as evolutionary/reinforcement learning agents. In some cases, they have achieved state-of-the-art results in open-ended environments like Minecraft Wang et al. [2023] or even improved mathematical conjectures, as in Romera-Paredes et al. [2024]. For example, in Wang et al. [2023], an LLM was used as an agent, with the state of the environment provided in the form of a prompt. General directives for the agent were added to the prompt, along with an explicit request to propose an achievable task considering the current situation of the environment. An example of a general directive could be a request to explore the environment as much as possible or an example of a simple task in a situation. This task was then translated by another LLM into code as a possible action for the current state of the environment. In this way, by using a combination of novelty search and in-context learning, it is possible to leverage the world knowledge of the LLM to achieve state-of-the-art results, while being able to answer to a given situation/state with a non atomistic/granular action.

3 Vision

The idea of using LLMs as agents could also be beneficial in IFE research by addressing some of the issues faced by the approaches in Jonas Degrave [2022] and Seo et al. [2024].

Figure 3 depicts a high level view of the envisioned ICF system, where experimental, diagnostics and simulations data augmented by the background knowledge used by LLM and DRL models to guide the research process.

Our plan of investigation is as follows:

1. Retain the global structure of the methodologies in Jonas Degrave [2022] and Seo et al. [2024], but use an LLM to shape more effective rewards or sub-tactical goals. To achieve

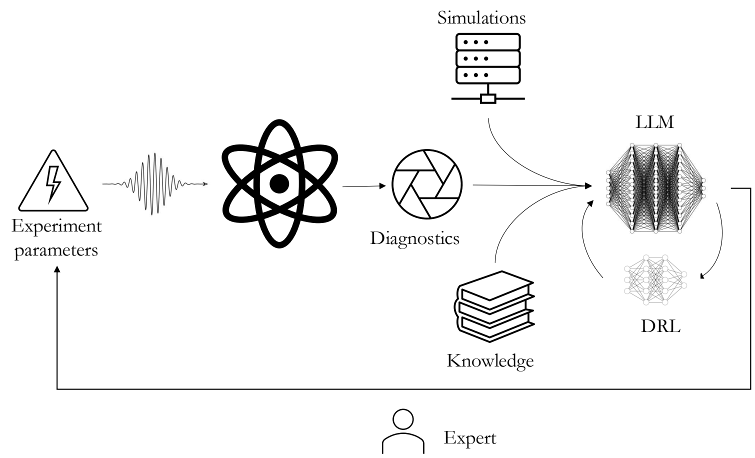

Figure 1: LLM-DRL guided IFE control loop.

this, we aim to use a known DRL architecture, like MPO and provide the problem at hand to the LLM in the form of a prompt with the aim to create a denser and more meaningful reward. The reward will be refined over several iterations. Similar ideas in different contexts have been explored in Kwon et al. [2023] and Xie et al. [2024].

2. Implement a Voyager Wang et al. [2023] approach to solve the environment at hand. In this way, we can avoid manually creating sub-tactical goals and instead use the LLM's knowledge to set them. This allows us to utilize rich state spaces without the need for a data-efficient algorithm, and additionally, devise actions as macro-tasks in the form of code, which overcomes a potential limitation of the approach in Seo et al. [2024].
3. Use two LLMs—one for reward shaping as described in point 1, and the other to act as an agent as described in point 2—thus combining the approaches outlined above.

4 Challenges

The use of LLMs in reinforcement learning in our context, comes with a number of challenges that need to be addressed.

Specifically, both Romera-Paredes et al. [2024] and Wang et al. [2023] utilize the world model of the LLM to solve their respective tasks. Although the improvement of the Collatz conjecture in Romera-Paredes et al. [2024] demonstrates the potential of LLMs in solving niche problems within human knowledge, this does not guarantee that another niche problem, such as inertial confinement nuclear fusion - particularly in the specific setting we aim to address — is adequately represented in the LLM. Therefore, it may be necessary to fine-tune the LLM before using it as an agent. This, of course, raises additional questions about the appropriate dataset for fine-tuning and which fine-tuning techniques should be employed, particularly in relation to the downstream agentic task. To that extent we plan to explore new science focused data sets and LLMs to be developed by the Frontiers in Artificial Intelligence for Science, Security and Technology (FASST) program of Crtitical and Emerging Technlogies [2024].

Additionally, the prompt schemes used in Romera-Paredes et al. [2024] and Wang et al. [2023] depend on an external operator with some knowledge of the problem being solved. Thus, while we may reduce the reliance on external knowledge, we do not entirely eliminate it. This only partially limits the scope of our approach, as it is likely that with minimal expert input, the LLM—through the feedback mechanisms described in Romera-Paredes et al. [2024] and Wang et al. [2023]—could generate sub-tactical goals or rewards that are superior to those conceived by a human.

We also need to mention challenges associated with the infrastructure requirements in support of our vision. Simulations, experiments and LLMs reside on different systems spanning from the edge to the cloud and supercomputers, thus solutions for data management, resources orchestration underpinned by appropriate security measures are very important for achieving our goals.

5 Implementation

Implementation of the presented vision can be divided into two major categories: 1. LLM-DRL pipeline development, including LLM selection and fine tuning. 2. Data management and control loop development, including required infrastructure.

LLM-DRL pipeline development will include design of a series of templates based on expert knowledge. These templates will be provided to the LLM, which will generate outputs aimed at solving our specific task—in this case, controlling laser beams parameters to initiate nuclear fusion reaction. The generated output will be used to run a simulation, with the final state of the simulation, along with the output itself, serving as a new input for the LLM. If experts in the field can define clear sub-tactical goals for the injection control problem, we will combine the previous approach with a hierarchical method as described in Prakash et al. [2023]. This hierarchical approach complements the code-based output design, allowing us to create a library of codes for the laser beams, each corresponding to a specific sub-tactical goal. Additionally, we plan to use the LLM for reward shaping, as outlined in Kwon et al. [2023]. Reward shaping can be applied in the plain LLM-DRL approach, the hierarchical LLM-DRL approach, and in conjunction with other established methods. Finally, we will generate a series of high-quality trajectories using the LLM approach to apply an offline RL algorithm, such as behavior cloning Peng et al. [2018], Osa et al. [2018], Ho and Ermon [2016]. This will distill the knowledge of the LLM agent into a smaller neural network, better suited to meet the latency requirements of nuclear fusion devices.

LLMs, however, suffer from a series of limitations, the most significant of which is undoubtedly the presence of hallucinations. To mitigate hallucinations produced by LLMs, we will primarily use code-based outputs. Writing responses in code format has been shown to effectively reduce hallucinations Romera-Paredes et al. [2024], Wang et al. [2023]. Additionally, we aim to explore other approaches, such as using multiple LLMs in a sequential critique process, where each LLM evaluates and critiques the output of the previous one, or employing an ensemble of LLMs to provide a unified, consensus-based answer Chen et al. [2023], Shridhar et al. [2023], Madaan et al. [2023], Mousavi et al. [2023]. It is important to note that this approach can be enhanced with prompt tuning. For example, each LLM can be prompted to evaluate its own output. This self-assessment can be integrated into the ensemble or judging methodology we described. For instance, during the ensembling process, a weighted average could be applied based on the quality score each LLM assigns to its own output Zhang et al. [2024].

Regarding LLM selection and fine tuning, we will investigate a FUSION-LLM approach suggested in the recently published Chen et al. [2024], which is based on LLama Touvron et al. [2023] models. We will also consider alternatives like Granite Mishra et al. [2024] family of models and potential models coming from the FASST program mentioned previously. It is worth noting that fine-tuning, in our context, could involve collecting data pairs of diagnostic and laser signals from experiments, formatting them according to the prompt structure we plan to use for our LLM agents, and then initiating the fine-tuning process. This approach may be challenging, as our task could require tens of thousands of such data pairs. However, this requirement might be reduced if we can obtain high-quality data samples. We can augment experimental data with data from simulations, although the latter comes with the challenge of simulations needed to generate high fidelity data may be computationally expensive. An alternative approach is to use this data to train a reward signal and employ a Reinforcement Learning from Human Feedback (RLHF) framework, which typically requires fewer training samples to achieve optimal results Ouyang et al. [2022] Kaufmann et al. [2024].

Development of the control loop requires appropriate infrastructure to support orchestration of various components to ensure proper data management and support of LLM-DRL pipeline. Envisioned software stack will include combination of relevant high performance computing (HPC) and cloud native technologies. Various application programming interfaces (API) need to be created to connect different components of the control loop.

Last, but not least the implementation must provide safeguards to ensure that decisions by the LLM-DRL do not lead to catastrophic results. We must ensure that these decisions suggested by the models lie within acceptable boundaries, are explainable and verifiable by human experts. Care must be taken of the quality and reliability of data used for training and inference. Secure development and operation of the control system must be ensured at any given time across all components of the software and infrastructure.

6 Conclusions

Recent breakthrough in ICF advanced the science and confirmed the potential use of fusion as a source of abundant and clean energy. AI can play a crucial role in developing fusion technology further. We presented our vision of using LLM and DRL for guiding IFE research. We discussed possible implementation of our the proposed approach both from the algorithmic and infrastructure point of view.

To conclude we want to emphasize that our vision presumes collaboration between scientists and engineers from different domains and multiple organizations. Specifically, we encourage the machine learning community to engage in this work through physics-informed approaches and neural operators, enabling the use of surrogate models to simplify and accelerate complex simulations. We also see this work as a potential source of inspiration for the LLM and DRL communities. It could serve as a novel testing ground for advanced LLM agent algorithms, particularly in mathematical reasoning and planning. Additionally, the task's strict latency constraints make it an excellent benchmark for evaluating state-of-the-art distillation techniques in LLM planning systems.

We are actively engaged with the IFE community through discussions with STFC, University of Oxford and Department of Energy. These organizations involved in high energy density research can contribute by providing experimental and simulations data, validating capabilities developed through our proposed approach and adopting them for production IFE workflows.

References

H. Abu-Shawareb, R. Acree, P. Adams, J. Adams, B. Addis, R. Aden, P. Adrian, B. B. Afeyan, M. Aggleton, L. Aghaian, A. Aguirre, D. Aikens, J. Akre, F. Albert, M. Albrecht, B. J. Albright, J. Albritton, J. Alcala, C. Alday, D. A. Alessi, N. Alexander, J. Alfonso, N. Alfonso, E. Alger, S. J. Ali, Z. A. Ali, W. E. Alley, P. Amala, P. A. Amendt, P. Amick, S. Ammula, C. Amorin, D. J. Ampleford, R. W. Anderson, T. Anklam, N. Antipa, B. Appelbe, C. Aracne-Ruddle, E. Araya, M. Arend, P. Arnold, T. Arnold, J. Asay, L. J. Atherton, D. Atkinson, R. Atkinson, J. M. Auerbach, B. Austin, L. Auyang, A. S. Awwal, J. Ayers, S. Ayers, T. Ayers, S. Azevedo, B. Bachmann, C. A. Back, J. Bae, D. S. Bailey, J. Bailey, T. Baisden, K. L. Baker, H. Baldis, D. Barber, M. Barberis, D. Barker, A. Barnes, C. W. Barnes, M. A. Barrios, C. Barty, I. Bass, S. H. Batha, S. H. Baxamusa, G. Bazan, J. K. Beagle, R. Beale, B. R. Beck, J. B. Beck, M. Bedzyk, R. G. Beeler, R. G. Beeler, W. Behrendt, L. Belk, P. Bell, M. Belyaev, J. F. Benage, G. Bennett, L. R. Benedetti, L. X. Benedict, R. Berger, T. Bernat, L. A. Bernstein, B. Berry, L. Bertolini, G. Besenbruch, J. Betcher, R. Bettenhausen, R. Betti, B. Bezzerides, S. D. Bhandarkar, R. Bickel, J. Biener, T. Biesiada, K. Bigelow, J. Bigelow-Granillo, V. Bigman, R. M. Bionta, N. W. Birge, M. Bitter, A. C. Black, R. Bleile, D. L. Bleuel, E. Bliss, E. Bliss, B. Blue, T. Boehly, K. Boehm, C. D. Boley, R. Bonanno, E. J. Bond, T. Bond, M. J. Bonino, M. Borden, J.-L. Bourgade, J. Bousquet, J. Bowers, M. Bowers, R. Boyd, A. Bozek, D. K. Bradley, K. S. Bradley, P. A. Bradley, L. Bradley, L. Brannon, P. S. Brantley, D. Braun, T. Braun, K. Brienza-Larsen, T. M. Briggs, J. Britten, E. D. Brooks, D. Browning, M. W. Bruhn, T. A. Brunner, H. Bruns, G. Brunton, B. Bryant, T. Buczek, J. Bude, L. Buitano, S. Burkhart, J. Burmark, A. Burnham, R. Burr, L. E. Busby, B. Butlin, R. Cabeltis, M. Cable, W. H. Cabot, B. Cagadas, J. Caggiano, R. Cahayag, S. E. Caldwell, S. Calkins, D. A. Callahan, J. Calleja-Aguirre, L. Camara, D. Camp, E. M. Campbell, J. H. Campbell, B. Carey, R. Carey, K. Carlisle, L. Carlson, L. Carman, J. Carmichael, A. Carpenter, C. Carr, J. A. Carrera, D. Casavant, A. Casey, D. T. Casey, A. Castillo, E. Castillo, J. I. Castor, C. Castro, W. Caughey, R. Cavitt, J. Celeste, P. M. Celliers, C. Cerjan, G. Chandler, B. Chang, C. Chang, J. Chang, L. Chang, R. Chapman, T. Chapman, L. Chase, H. Chen, H. Chen, K. Chen, L.-Y. Chen, B. Cheng, J. Chittenden, C. Choate, J. Chou, R. E. Chrien, M. Chrisp, K. Christensen, M. Christensen, A. R. Christopherson, M. Chung, J. A. Church, A. Clark, D. S. Clark, K. Clark, R. Clark, L. Claus,

B. Cline, J. A. Cline, J. A. Cobble, K. Cochrane, B. Cohen, S. Cohen, M. R. Collette, G. Collins, L. A. Collins, T. J. B. Collins, A. Conder, B. Conrad, M. Conyers, A. W. Cook, D. Cook, R. Cook, J. C. Cooley, G. Cooper, T. Cope, S. R. Copeland, F. Coppari, J. Cortez, J. Cox, D. H. Crandall, J. Crane, R. S. Craxton, M. Cray, A. Crilly, J. W. Crippen, D. Cross, M. Cuneo, G. Cuotts, C. E. Czajka, D. Czechowicz, T. Daly, P. Danforth, R. Darbee, B. Darlington, P. Datte, L. Dauffy, G. Davalos, S. Davidovits, P. Davis, J. Davis, S. Dawson, R. D. Day, T. H. Day, M. Dayton, C. Deck, C. Decker, C. Deeney, K. A. DeFriend, G. Deis, N. D. Delamater, J. A. Delettrez, R. Demaret, S. Demos, S. M. Dempsey, R. Desjardin, T. Desjardins, M. P. Desjarlais, E. L. Dewald, J. DeYoreo, S. Diaz, G. Dimonte, T. R. Dittrich, L. Divol, S. N. Dixit, J. Dixon, E. S. Dodd, D. Dolan, A. Donovan, M. Donovan, T. Döppner, C. Dorrer, N. Dorsano, M. R. Douglas, D. Dow, J. Downie, E. Downing, M. Dozieres, V. Draggoo, D. Drake, R. P. Drake, T. Drake, G. Dreifuerst, D. F. DuBois, P. F. DuBois, G. Dunham, R. Dylla-Spears, A. K. L. Dymoke-Bradshaw, B. Dzenitis, C. Ebbers, M. Eckart, S. Eddinger, D. Eder, D. Edgell, M. J. Edwards, P. Efthimion, J. H. Eggert, B. Ehrlich, P. Ehrmann, S. Elhadj, C. Ellerbee, N. S. Elliott, C. L. Ellison, F. Elsner, M. Emerich, K. Engelhorn, T. England, E. English, P. Epperson, R. Epstein, G. Erbert, M. A. Erickson, D. J. Erskine, A. Erlandson, R. J. Espinosa, C. Estes, K. G. Estabrook, S. Evans, A. Fabyan, J. Fair, R. Fallejo, N. Farmer, W. A. Farmer, M. Farrell, V. E. Fatherley, M. Fedorov, E. Feigenbaum, M. Feit, W. Ferguson, J. C. Fernandez, A. Fernandez-Panella, S. Fess, J. E. Field, C. V. Filip, J. R. Fincke, T. Finn, S. M. Finnegan, R. G. Finucane, M. Fischer, A. Fisher, J. Fisher, B. Fishler, D. Fittinghoff, P. Fitzsimmons, M. Flegel, K. A. Flippo, J. Florio, J. Folta, P. Folta, L. R. Foreman, C. Forrest, A. Forsman, J. Fooks, M. Foord, R. Fortner, K. Fournier, D. E. Fratanduono, N. Frazier, T. Frazier, C. Frederick, M. S. Freeman, J. Frenje, D. Frey, G. Frieders, S. Friedrich, D. H. Froula, J. Fry, T. Fuller, J. Gaffney, S. Gales, B. Le Galloudec, K. K. Le Galloudec, A. Gambhir, L. Gao, W. J. Garbett, A. Garcia, C. Gates, E. Gaut, P. Gauthier, Z. Gavin, J. Gaylord, M. Geissel, F. Génin, J. Georgeson, H. Geppert-Kleinrath, V. Geppert-Kleinrath, N. Gharibyan, J. Gibson, C. Gibson, E. Giraldez, V. Glebov, S. G. Glendinning, S. Glenn, S. H. Glenzer, S. Goade, P. L. Gobby, S. R. Goldman, B. Golick, M. Gomez, V. Goncharov, D. Goodin, P. Grabowski, E. Grafil, P. Graham, J. Grandy, E. Grasz, F. Graziani, G. Greenman, J. A. Greenough, A. Greenwood, G. Gregori, T. Green, J. R. Griego, G. P. Grim, J. Grondalski, S. Gross, J. Guckian, N. Guler, B. Gunney, G. Guss, S. Haan, J. Hackbarth, L. Hackel, R. Hackel, C. Haefner, C. Hagmann, K. D. Hahn, S. Hahn, B. J. Haid, B. M. Haines, B. M. Hall, C. Hall, G. N. Hall, M. Hamamoto, S. Hamel, C. E. Hamilton, B. A. Hammel, J. H. Hammer, G. Hampton, A. Hamza, A. Handler, S. Hansen, D. Hanson, R. Haque, D. Harding, E. Harding, J. D. Hares, D. B. Harris, J. A. Harte, E. P. Hartouni, R. Hatarik, S. Hatchett, A. A. Hauer, M. Havre, R. Hawley, J. Hayes, J. Hayes, S. Hayes, A. Hayes-Sterbenz, C. A. Haynam, D. A. Haynes, D. Headley, A. Heal, J. E. Heebner, S. Heerey, G. M. Heestand, R. Heeter, N. Hein, C. Heinbockel, C. Hendricks, M. Henesian, J. Heninger, J. Henrikson, E. A. Henry, E. B. Herbold, M. R. Hermann, G. Hermes, J. E. Hernandez, V. J. Hernandez, M. C. Herrmann, H. W. Herrmann, O. D. Herrera, D. Hewett, R. Hibbard, D. G. Hicks, D. Hill, K. Hill, T. Hilsabeck, D. E. Hinkel, D. D. Ho, V. K. Ho, J. K. Hoffer, N. M. Hoffman, M. Hohenberger, M. Hohensee, W. Hoke, D. Holdener, F. Holdener, J. P. Holder, B. Holko, D. Holunga, J. F. Holzrichter, J. Honig, D. Hoover, D. Hopkins, L. Berzak Hopkins, M. Hoppe, M. L. Hoppe, J. Horner, R. Hornung, C. J. Horsfield, J. Horvath, D. Hotaling, R. House, L. Howell, W. W. Hsing, S. X. Hu, H. Huang, J. Huckins, H. Hui, K. D. Humbird, J. Hund, J. Hunt, O. A. Hurricane, M. Hutton, K. H.-K. Huynh, L. Inandan, C. Iglesias, I. V. Igumenshchev, N. Izumi, M. Jackson, J. Jackson, S. D. Jacobs, G. James, K. Jancaitis, J. Jarboe, L. C. Jarrott, D. Jasion, J. Jaquez, J. Jeet, A. E. Jenei, J. Jensen, J. Jimenez, R. Jimenez, D. Jobe, Z. Johal, H. M. Johns, D. Johnson, M. A. Johnson, M. Gatu Johnson, R. J. Johnson, S. Johnson, S. A. Johnson, T. Johnson, K. Jones, O. Jones, M. Jones, R. Jorge, H. J. Jorgenson, M. Julian, B. I. Jun, R. Jungquist, J. Kaae, N. Kabadi, D. Kaczala, D. Kalantar, K. Kangas, V. V. Karasiev, M. Karasik, V. Karpenko, A. Kasarky, K. Kasper, R. Kauffman, M. I. Kaufman, C. Keane, L. Keaty, L. Kegelmeyer, P. A. Keiter, P. A. Kellett, J. Kellogg, J. H. Kelly, S. Kemic, A. J. Kemp, G. E. Kemp, G. D. Kerbel, D. Kershaw, S. M. Kerr, T. J. Kessler, M. H. Key, S. F. Khan, H. Khater, C. Kiikka, J. Kilkenny, Y. Kim, Y.-J. Kim, J. Kimko, M. Kimmel, J. M. Kindel, J. King, R. K. Kirkwood, L. Klaus, D. Klem, J. L. Kline, J. Klingmann, G. Kluth, P. Knapp, J. Knauer, J. Knipping, M. Knudson, D. Kobs, J. Koch, T. Kohut, C. Kong, J. M. Koning, P. Koning, S. Konior, H. Kornblum, L. B. Kot, B. Kozioziemski, M. Kozlowski, P. M. Kozlowski, J. Krammen, N. S. Krasheninnikova, B. Kraus, W. Krauser, J. D. Kress, A. L. Kritcher, E. Krieger, J. J. Kroll, W. L. Kruer, M. K. G. Kruse, S. Kucheyev, M. Kumbera, S. Kumpan, J. Kunimune, B. Kustowski, T. J. T. Kwan, G. A. Kyrala, S. Laffite, M. Lafon, K. LaFortune, B. Lahmann, B. Lairson, O. L. Landen, J. Langenbrunner, L. Lagin,

T. Land, M. Lane, D. Laney, A. B. Langdon, S. H. Langer, A. Langro, N. E. Lanier, T. E. Lanier, D. Larson, B. F. Lasinski, D. Lassle, D. LaTray, G. Lau, N. Lau, C. Laumann, A. Laurence, T. A. Laurence, J. Lawson, H. P. Le, R. R. Leach, L. Leal, A. Leatherland, K. LeChien, B. Lechleiter, A. Lee, M. Lee, T. Lee, R. J. Leeper, E. Lefebvre, J.-P. Leidinger, B. LeMire, R. W. Lemke, N. C. Lemos, S. Le Pape, R. Lerche, S. Lerner, S. Letts, K. Levedahl, T. Lewis, C. K. Li, H. Li, J. Li, W. Liao, Z. M. Liao, D. Liedahl, J. Liebman, G. Lindford, E. L. Lindman, J. D. Lindl, H. Loey, R. A. London, F. Long, E. N. Loomis, F. E. Lopez, H. Lopez, E. Losbanos, S. Loucks, R. Lowe-Webb, E. Lundgren, A. P. Ludwigsen, R. Luo, J. Lusk, R. Lyons, T. Ma, Y. Macallop, M. J. MacDonald, B. J. MacGowan, J. M. Mack, A. J. Mackinnon, S. A. MacLaren, A. G. MacPhee, G. R. Magelssen, J. Magoon, R. M. Malone, T. Malsbury, R. Managan, R. Mancini, K. Manes, D. Maney, D. Manha, O. M. Mannion, A. M. Manuel, E. Mapoles, G. Mara, T. Marcotte, E. Marin, M. M. Marinak, C. Mariscal, D. A. Mariscal, E. F. Mariscal, E. V. Marley, J. A. Marozas, R. Marquez, C. D. Marshall, F. J. Marshall, M. Marshall, S. Marshall, J. Marticorena, D. Martinez, I. Maslennikov, D. Mason, R. J. Mason, L. Masse, W. Massey, P.-E. Masson-Laborde, N. D. Masters, D. Mathisen, E. Mathison, J. Matone, M. J. Matthews, C. Mattoon, T. R. Mattsson, K. Matzen, C. W. Mauche, M. Mauldin, T. McAbee, M. McBurney, T. Mccarville, R. L. McCrory, A. M. McEvoy, C. McGuffey, M. Mcinnis, P. McKenty, M. S. McKinley, J. B. McLeod, A. McPherson, B. Mcquillan, M. Meamber, K. D. Meaney, N. B. Meezan, R. Meissner, T. A. Mehlhorn, N. C. Mehta, J. Menapace, F. E. Merrill, B. T. Merritt, E. C. Merritt, D. D. Meyerhofer, S. Mezyk, R. J. Mich, P. A. Michel, D. Milam, C. Miller, D. Miller, D. S. Miller, E. Miller, E. K. Miller, J. Miller, M. Miller, P. E. Miller, T. Miller, W. Miller, V. Miller-Kamm, M. Millot, J. L. Milovich, P. Minner, J.-L. Miquel, S. Mitchell, K. Molvig, R. C. Montesanti, D. S. Montgomery, M. Monticelli, A. Montoya, J. D. Moody, A. S. Moore, E. Moore, M. Moran, J. C. Moreno, K. Moreno, B. E. Morgan, T. Morrow, J. W. Morton, E. Moses, K. Moy, R. Muir, M. S. Murillo, J. E. Murray, J. R. Murray, D. H. Munro, T. J. Murphy, F. M. Munteanu, J. Nafziger, T. Nagayama, S. R. Nagel, R. Nast, R. A. Negres, A. Nelson, D. Nelson, J. Nelson, S. Nelson, S. Nemethy, P. Neumayer, K. Newman, M. Newton, H. Nguyen, J.-M. G. Di Nicola, P. Di Nicola, C. Niemann, A. Nikroo, P. M. Nilson, A. Nobile, V. Noorai, R. Nora, M. Norton, M. Nostrand, V. Note, S. Novell, P. F. Nowak, A. Nunez, R. A. Nyholm, M. O'Brien, A. Oceguera, J. A. Oertel, J. Okui, B. Olejniczak, J. Oliveira, P. Olsen, B. Olson, K. Olson, R. E. Olson, Y. P. Opachich, N. Orsi, C. D. Orth, M. Owen, S. Padalino, E. Padilla, R. Paguio, S. Paguio, J. Paisner, S. Pajoom, A. Pak, S. Palaniyappan, K. Palma, T. Pannell, F. Papp, D. Paras, T. Parham, H.-S. Park, A. Pasternak, S. Patankar, M. V. Patel, P. K. Patel, R. Patterson, S. Patterson, B. Paul, M. Paul, E. Pauli, O. T. Pearce, J. Pearcy, B. Pedrotti, A. Peer, L. J. Pelz, B. Penetrante, J. Penner, A. Perez, L. J. Perkins, E. Pernice, T. S. Perry, S. Person, D. Petersen, T. Petersen, D. L. Peterson, E. B. Peterson, J. E. Peterson, J. L. Peterson, K. Peterson, R. R. Peterson, R. D. Petrasso, F. Philippe, T. J. Phipps, E. Piceno, Y. Ping, L. Pickworth, J. Pino, R. Plummer, G. D. Pollack, S. M. Pollaine, B. B. Pollock, D. Ponce, J. Ponce, J. Pontelandolfo, J. L. Porter, J. Post, O. Poujade, C. Powell, H. Powell, G. Power, M. Pozulp, M. Prantil, M. Prasad, S. Pratuch, S. Price, K. Primdahl, S. Prisbrey, R. Procassini, A. Pruyne, B. Pudliner, S. R. Qiu, K. Quan, M. Quinn, J. Quintenz, P. B. Radha, F. Rainer, J. E. Ralph, K. S. Raman, R. Raman, P. Rambo, S. Rana, A. Randewich, D. Rardin, M. Ratledge, N. Ravelo, F. Ravizza, M. Rayce, A. Raymond, B. Raymond, B. Reed, C. Reed, S. Regan, B. Reichelt, V. Reis, S. Reisdorf, V. Rekow, B. A. Remington, A. Rendon, W. Requieron, M. Rever, H. Reynolds, J. Reynolds, J. Rhodes, M. Rhodes, M. C. Richardson, B. Rice, N. G. Rice, R. Rieben, A. Rigatti, S. Riggs, H. G. Rinderknecht, K. Ring, B. Riordan, R. Riquier, C. Rivers, D. Roberts, V. Roberts, G. Robertson, H. F. Robey, J. Robles, P. Rocha, G. Rochau, J. Rodriguez, S. Rodriguez, M. Rosen, M. Rosenberg, G. Ross, J. S. Ross, P. Ross, J. Rouse, D. Rovang, A. M. Rubenchik, M. S. Rubery, C. L. Ruiz, M. Rushford, B. Russ, J. R. Rygg, B. S. Ryujin, R. A. Sacks, R. F. Sacks, K. Saito, T. Salmon, J. D. Salmonson, J. Sanchez, S. Samuelson, M. Sanchez, C. Sangster, A. Saroyan, J. Sater, A. Satsangi, S. Sauers, R. Saunders, J. P. Sauppe, R. Sawicki, D. Sayre, M. Scanlan, K. Schaffers, G. T. Schappert, S. Schiaffino, D. J. Schlossberg, D. W. Schmidt, M. J. Schmitt, D. H. G. Schneider, M. B. Schneider, R. Schneider, M. Schoff, M. Schollmeier, M. Schölmerich, C. R. Schroeder, S. E. Schrauth, H. A. Scott, I. Scott, J. M. Scott, R. H. H. Scott, C. R. Scullard, T. Sedillo, F. H. Seguin, W. Seka, J. Senecal, S. M. Sepke, L. Seppala, K. Sequoia, J. Severyn, J. M. Sevier, N. Sewell, S. Seznec, R. C. Shah, J. Shamlian, D. Shaughnessy, M. Shaw, R. Shaw, C. Shearer, R. Shelton, N. Shen, M. W. Sherlock, A. I. Shestakov, E. L. Shi, S. J. Shin, N. Shingleton, W. Shmayda, M. Shor, M. Shoup, C. Shuldberg, L. Siegel, F. J. Silva, A. N. Simakov, B. T. Sims, D. Sinars, P. Singh, H. Sio, K. Skulina, S. Skupsky, S. Slutz, M. Sluyter, V. A. Smalyuk, D. Smauley, R. M. Smeltser, C. Smith, I. Smith, J. Smith, L. Smith, R. Smith,

R. Sohn, S. Sommer, C. Sorce, M. Sorem, J. M. Soures, M. L. Spaeth, B. K. Spears, S. Speas, D. Speck, R. Speck, J. Spears, T. Spinka, P. T. Springer, M. Stadermann, B. Stahl, J. Stahoviak, L. G. Stanton, R. Steele, W. Steele, D. Steinman, R. Stemke, R. Stephens, S. Sterbenz, P. Sterne, D. Stevens, J. Stevers, C. B. Still, C. Stoeckl, W. Stoeffl, J. S. Stolken, C. Stolz, E. Storm, G. Stone, S. Stoupin, E. Stout, I. Stowers, R. Strauser, H. Streckart, J. Streit, D. J. Strozzi, T. Suratwala, G. Sutcliffe, L. J. Suter, S. B. Sutton, V. Svidzinski, G. Swadling, W. Sweet, A. Szoke, M. Tabak, M. Takagi, A. Tambazidis, V. Tang, M. Taranowski, L. A. Taylor, S. Telford, W. Theobald, M. Thi, A. Thomas, C. A. Thomas, I. Thomas, R. Thomas, I. J. Thompson, A. Thongstisubskul, C. B. Thorsness, G. Tietbohl, R. E. Tipton, M. Tobin, N. Tomlin, R. Tommasini, A. J. Toreja, J. Torres, R. P. J. Town, S. Townsend, J. Trenholme, A. Trivelpiece, C. Trosseille, H. Truax, D. Trummer, S. Trummer, T. Truong, D. Tubbs, E. R. Tubman, T. Tunnell, D. Turnbull, R. E. Turner, M. Ulitsky, R. Upadhye, J. L. Vaher, P. VanArsdall, D. VanBlarcom, M. Vandenboomgaerde, R. VanQuinlan, B. M. Van Wonterghem, W. S. Varnum, A. L. Velikovich, A. Vella, C. P. Verdon, B. Vermillion, S. Vernon, R. Vesey, J. Vickers, R. M. Vignes, M. Visosky, J. Vocke, P. L. Volegov, S. Vonhof, R. Von Rotz, H. X. Vu, M. Vu, D. Wall, J. Wall, R. Wallace, B. Wallin, D. Walmer, C. A. Walsh, C. F. Walters, C. Waltz, A. Wan, A. Wang, Y. Wang, J. S. Wark, B. E. Warner, J. Watson, R. G. Watt, P. Watts, J. Weaver, R. P. Weaver, S. Weaver, C. R. Weber, P. Weber, S. V. Weber, P. Wegner, B. Welday, L. Welser-Sherrill, K. Weiss, K. Widmann, G. F. Wheeler, W. Whistler, R. K. White, H. D. Whitley, P. Whitman, M. E. Wickett, C. Widmayer, J. Wiedwald, R. Wilcox, S. Wilcox, C. Wild, B. H. Wilde, C. H. Wilde, K. Wilhelmsen, M. D. Wilke, H. Wilkens, P. Wilkins, S. C. Wilks, E. A. Williams, G. J. Williams, W. Williams, W. H. Williams, D. C. Wilson, B. Wilson, E. Wilson, R. Wilson, S. Winters, J. Wisoff, M. Wittman, J. Wolfe, A. Wong, K. W. Wong, L. Wong, N. Wong, R. Wood, D. Woodhouse, J. Woodruff, D. T. Woods, S. Woods, B. N. Woodworth, E. Wooten, A. Wootton, K. Work, J. B. Workman, J. Wright, M. Wu, C. Wuest, F. J. Wysocki, H. Xu, M. Yamaguchi, B. Yang, S. T. Yang, J. Yatabe, C. B. Yeamans, B. C. Yee, S. A. Yi, L. Yin, B. Young, C. S. Young, C. V. Young, P. Young, K. Youngblood, R. Zacharias, G. Zagaris, N. Zaitseva, F. Zaka, F. Ze, B. Zeiger, M. Zika, G. B. Zimmerman, T. Zobrist, J. D. Zuegel, and A. B. Zylstra. Lawson criterion for ignition exceeded in an inertial fusion experiment. *Phys. Rev. Lett.*, 129:075001, Aug 2022. doi: 10.1103/PhysRevLett.129.075001. URL `https://link.aps.org/doi/10.1103/PhysRevLett.129.075001`.

Francis F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer Cham, 3 edition, 2016.

Huaben Chen, Wenkang Ji, Lufeng Xu, and Shiyu Zhao. Multi-agent consensus seeking via large language models. *arXiv preprint arXiv:2310.20151*, 2023. URL `https://westlakeintelligentrobotics.github.io/ConsensusLLM/`. arXiv preprint arXiv:2310.20151.

Mingkai Chen, Taowen Wang, James Chenhao Liang, Chuan Liu, Chunshu Wu, Qifan Wang, Ying Nian Wu, Michael Huang, Chuang Ren, Ang Li, Tong Geng, and Dongfang Liu. Inertial confinement fusion forecasting via llms, 2024. URL `https://arxiv.org/abs/2407.11098`.

DOE. Doe explains...tokamaks, 2023. URL `https://link.aps.org/doi/10.1103/PhysRevLett.129.075001`.

National Ignition Facility. Achieving fusion ignition, 2024. URL `https://lasers.llnl.gov/science/achieving-fusion-ignition`.

V. Gopalaswamy, R. Betti, J. P. Knauer, N. Luciani, D. Patel, K. M. Woo, A. Bose, I. V. Igumenshchev, E. M. Campbell, K. S. Anderson, K. A. Bauer, M. J. Bonino, D. Cao, A. R. Christopherson, G. W. Collins, T. J. B. Collins, J. R. Davies, J. A. Delettrez, D. H. Edgell, R. Epstein, C. J. Forrest, D. H. Froula, V. Y. Glebov, V. N. Goncharov, D. R. Harding, S. X. Hu, Jacobs, R. T. Janezic, J. H. Kelly, O. M. Mannion, A. Maximov, F. J. Marshall, D. T. Michel, S. Miller, S. F. B. Morse, J. Palastro, J. Peebles, P. B. Radha, S. P. Regan, S. Sampat, T. C. Sangster, A. B. Sefkow, W. Seka, R. C. Shah, W. T. Shmyada, A. Shvydky, C. Stoeckl, A. A. Solodov, W. Theobald, J. D. Zuegel, M. Gatu Johnson, R. D. Petrasso, C. K. Li, and J. A. Frenje. Tripled yield in direct-drive laser fusion through statistical modelling. *Nature*, 565:581–586, 2019.

Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Proceedings of the 30th Conference on Neural Information Processing Systems (NeurIPS*

2016), Barcelona, Spain, 2016. NeurIPS. URL https://papers.nips.cc/paper/6391-generative-adversarial-imitation-learning.

F. Hofmann, J. B. Lister, M. Anton, S. Barry, R. Behn, S. Bernel, G. Besson, F. Bühlmann, R. Chavan, M. Corboz, M. J. Dutch, B. P. Duval, D. Fasel, A. Favre, S. Franke, A. Heym, A. Hirt, Ch. Hollenstein, P. Isoz, B. Joye, X. Llobet, J. C. Magnin, B. Marletaz, Ph. Marmillod, Y. Martin, J. M. Mayor, J. M. Moret, Ch. Nieswand, P. J. Paris, A. Perez, Z. A. Pietrzyk, R. A. Pitts, A. Pochelon, R. Rage, O. Sauter, G. Tonetti, M. Q. Tran, F. Troyon, D. J. Ward, and H. Weisen. Creation and control of variably shaped plasmas in tcv. *Plasma Physics and Controlled Fusion*, 36:B277–B287, 1994. doi: 10.1088/0741-3335/36/12B/023. URL https://doi.org/10.1088/0741-3335/36/12B/023.

Jonas Buchli et al. Jonas Degrave, Federico Felici. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602:414–419, February 2022. doi: 10.1038/s41586-021-04301-9. URL https://doi.org/10.1038/s41586-021-04301-9.

Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. A survey of reinforcement learning from human feedback. *arXiv preprint arXiv:2312.14925*, 2024. URL https://arxiv.org/abs/2312.14925. arXiv preprint arXiv:2312.14925.

Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. *arXiv preprint arXiv:2303.00001*, 2023. URL https://arxiv.org/abs/2303.00001. Published as a conference paper at ICLR 2023.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023. URL https://selfrefine.info/. arXiv preprint arXiv:2303.17651.

Mayank Mishra, Matt Stallone, Gaoyuan Zhang, Yikang Shen, Aditya Prasad, Adriana Meza Soria, Michele Merler, Parameswaran Selvam, Saptha Surendran, Shivdeep Singh, Manish Sethi, Xuan-Hong Dang, Pengyuan Li, Kun-Lung Wu, Syed Zawad, Andrew Coleman, Matthew White, Mark Lewis, Raju Pavuluri, Yan Koyfman, Boris Lublinsky, Maximilien de Bayser, Ibrahim Abdelaziz, Kinjal Basu, Mayank Agarwal, Yi Zhou, Chris Johnson, Aanchal Goyal, Hima Patel, Yousaf Shah, Petros Zerfos, Heiko Ludwig, Asim Munawar, Maxwell Crouse, Pavan Kapanipathi, Shweta Salaria, Bob Calio, Sophia Wen, Seetharami Seelam, Brian Belgodere, Carlos Fonseca, Amith Singhee, Nirmit Desai, David D. Cox, Ruchir Puri, and Rameswar Panda. Granite code models: A family of open foundation models for code intelligence, 2024. URL https://arxiv.org/abs/2405.04324.

Sajad Mousavi, Ricardo Luna Gutierrez, Desik Rengarajan, Vineet Gundecha, Ashwin Ramesh Babu, Avisek Naug, Antonio Guillen, and Soumyendu Sarkar. N-critics: Self-refinement of large language models with ensemble of critics. *arXiv preprint arXiv:2310.18679*, 2023. URL https://arxiv.org/abs/2310.18679. Presented at NeurIPS 2023 Workshop on Robustness of Few-shot and Zero-shot Learning in Foundation Models.

John Nuckolls, Lowell Wood, Albert Thiessen, and George Zimmerman. Laser compression of matter to super-high densities: Thermonuclear (ctr) applications. *Nature*, 239:139–142, 1972.

Office of Crtitical and U.S. Department of Eenrgy Emerging Technlogies. Frontiers in artificial intelligence for science, security and technology (fasst), 2024. URL https://www.energy.gov/fasst.

Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 7(1-2): 1–179, 2018. doi: 10.1561/2300000053. URL https://doi.org/10.1561/2300000053.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*, 2022. URL https://arxiv.org/abs/2203.02155. arXiv preprint arXiv:2203.02155.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)*, 37(4):143:1–143:14, August 2018. ISSN 0730-0301. doi: 10.1145/3197517.3201311. URL `https://doi.org/10.1145/3197517.3201311`.

Bharat Prakash, Tim Oates, and Tinoosh Mohsenin. Llm augmented hierarchical agents. *arXiv preprint arXiv:2311.05596*, 2023. URL `https://arxiv.org/abs/2311.05596`. Presented at the LangRob Workshop @ Conference on Robot Learning (CoRL 2023).

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. Mathematical discoveries from program search with large language models. *Nature*, 2024. doi: 10.1038/s41586-023-06924-6. URL `https://doi.org/10.1038/s41586-023-06924-6`.

Jaemin Seo, SangKyeun Kim, Azarakhsh Jalalvand, Rory Conlin, Andrew Rothstein, Joseph Abbate, Keith Erickson, Josiah Wai, Ricardo Shousha, and Egemen Kolemen. Avoiding fusion plasma tearing instability with deep reinforcement learning. *Nature*, 626:746–751, February 2024. doi: 10.1038/s41586-024-07024-9. URL `https://doi.org/10.1038/s41586-024-07024-9`.

Kumar Shridhar, Koustuv Sinha, Andrew Cohen, Tianlu Wang, Ping Yu, Ram Pasunuru, Mrinmaya Sachan, Jason Weston, and Asli Celikyilmaz. The art of llm refinement: Ask, refine, and trust. *arXiv preprint arXiv:2311.07961*, 2023. URL `https://arxiv.org/abs/2311.07961`. arXiv preprint arXiv:2311.07961.

Vladimir Tikhonchuk. *Particle Kinetics and Laser-Plasma Interactions*. Cambridge Scholars Publishing, 1 edition, 2013.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023. URL `https://arxiv.org/abs/2302.13971`.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi "Jim" Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291v2*, 2023. URL `https://arxiv.org/abs/2305.16291v2`.

Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanchao Yang, and Tao Yu. Text2reward: Reward shaping with language models for reinforcement learning. *arXiv preprint arXiv:2309.11489v3*, 2024. URL `https://arxiv.org/abs/2309.11489v3`. Published as a conference paper at ICLR 2024.

Xiaoying Zhang, Baolin Peng, Ye Tian, Jingyan Zhou, Lifeng Jin, Linfeng Song, Haitao Mi, and Helen Meng. Self-alignment for factuality: Mitigating hallucinations in llms via self-evaluation. *arXiv preprint arXiv:2402.09267*, 2024. URL `https://arxiv.org/abs/2402.09267`. arXiv preprint arXiv:2402.09267.