DAY ONE PROJECT

Focused Research Organizations to Accelerate Science, Technology, and Medicine

Samuel G. Rodriques
Adam H. Marblestone

September 2020

Summary

The next administration should rapidly create new Focused Research Organizations (FROs) to tackle scientific and technological challenges that cannot be efficiently addressed by standard organizational structures including academia, industry, National Laboratories, or Advanced Research Project Agencies (e.g., DARPA). FROs would be independent from existing universities or labs, focused on a single basic science or technology problem, and organized similarly to a startup. FROs would fill a key structural gap in our nation's research and development (R&D) system, enabling major advances in areas that (i) require levels of coordinated engineering or system-building inaccessible to academia, (ii) benefit society broadly in ways that industry cannot rapidly monetize, and (iii) harbor opportunities for acceleration through innovative new technologies and processes. Each FRO would produce a well-defined tool or technology, a key scientific dataset, or a refined process or resource that would dramatically boost progress and help maintain U.S. competitiveness in a broad technological or scientific field. Relevant areas for FROs include brain mapping, climate technology, biological tool and reagent development, data generation for preventative medicine, novel antibiotic development, nanofabrication, and more.

Challenge and Opportunity

The U.S. government is ill-equipped to fund R&D projects that require tight coordination and teamwork to create public goods. The majority of government-funded research outside of the defense sphere—including research funded through the National Institute of Health (NIH), the National Science Foundation (NSF), the Defense Advanced Research Projects Agency (DARPA), and the Advanced Research Projects Agency–Energy (ARPA-E)—is outsourced to externalized collaborations of university labs and/or commercial organizations. However, the academic reward structure favors individual credit and discourages systematic teamwork. Commercial incentives encourage teamwork but discourage the production of public goods. As a result, the United States is falling behind in key areas like microfabrication and human genomics to countries with greater abilities to centralize and accelerate focused research.

The solution is to enable the U.S. government to fund centralized research programs, termed Focused Research Organizations (FROs), to address well-defined challenges that require scale and coordination but that are not immediately profitable. FROs would be stand-alone "moonshot organizations" insulated from both academic and commercial incentive structures. FROs would be organized like startups, but they would pursue well-defined R&D goals in the public interest and would be accountable to their funding organizations rather than to shareholders. Each FRO would strive to accelerate a key R&D area via "multiplier effects" (such as dramatically reducing the cost of collecting critical scientific data), provide the United States with a decisive competitive advantage in that area, and de-risk substantial follow-on investment from the private and/or public sectors. Some FROs would lay the engineering foundations for subsequent government investment in programs similar in scope to the Human Genome Project.

Individual FRO-like entities have previously been established only occasionally and through disparate mechanisms. Most recently, the National Quantum Initiative Act established five FRO-like centers within National Labs, each funded at $25 million per year, to pursue advances in quantum-information technology. However, there is no systematic, agile process for the conception and creation of similar centers in a variety of fields. Establishing any FRO-like entity currently requires Congressional approval—an onerous and time-consuming process.

We expect FROs to attract broad bipartisan and popular support due to their potential to spawn new industries and establish American leadership. Precedent supports this expectation. The National Quantum Initiative Act, for instance, was co-sponsored by the bipartisan coalition of Lamar Alexander (R-TX), John Thune (R-SD), and Bill Nelson (D-FL), and passed the Senate by unanimous consent.

Plan of Action

The next administration should support the rapid establishment of 16 new FROs: four per year for the next four years, totaling 16 FROs. The next administration should work with Congress to secure new funding for these FROs, and the White House Office of Science and Technology Policy (OSTP) should oversee the development of a cross-disciplinary program to conceive and launch the FROs.

Funding

The total program budget for 16 FROs should be roughly $1 billion, or about $25–75 million per FRO allocated over 5-7 years (e.g., roughly $5–15 million per FRO per year). This is roughly 10 times the funding level accessible via a typical academic grant, yet comparable in cost to a DARPA project or to a philanthropic project like the Allen Institute's Mouse Brain Atlas (~$55 million). Moreover, this level of funding is similar to the funding needed by a Series A/B "hard tech" startup to achieve proof of concept for a new technology prior to commercialization. Funding should be authorized for the FRO program as whole rather than for each individual component. This will enable the program to move quickly and independently, in similar fashion to DARPA. Funding the program as a whole will also support cross-disciplinary FROs and FRO initiatives. Agencies such as NIH, NSF, the Department of Energy (DOE), the various ARPAs, or the "Directorate for Technology" proposed in the Endless Frontier Act could be involved in the FRO program and could solicit or put forward specific FROs.[1]

Logistics

FRO organization and operations should be designed to make FROs as agile, flexible, and self-directed as possible. Each FRO should exist independent of existing organizations such as National Laboratories or labs at other government agencies and academic institutions. Each FRO would be run by a CEO/CTO and staffed by a centralized, startup-like team of well-trained professionals sourced from both industry and academia. This personnel structure will enable

[1] 116th Congress. "S-3832 – Endless Frontier Act," (2020). https://www.congress.gov/bill/116th-congress/senate-bill/3832.

tighter alignment of team incentives and focus than would an externalized collaborative research program that uses existing entities (e.g., universities) as performers. This structure will also enable tighter alignment of incentives and focus than would a DARPA-like externalized effort coordinated by a single program manager (although some FROs could be created as an outcome or second stage of DARPA-like programs). Generally, FROs would rent commercial real estate for operations. In rare cases it may be appropriate for FROs to use National Lab facilities. Pay structure in FROs should be flexible to allow top talent to be recruited.

FROs should be expressly time-bound and outcome driven in order to prevent mission creep and organizational aging. This will require clear and pre-defined end-points/exits. As an FRO sunsets, stakeholders in that FRO's outputs should be convened to maximize output deployment and uptake. Intellectual property should be out-licensed or released publicly for similar reasons. Transition support should be provided to outgoing FRO employees. Follow-on from FROs could include formation and/or incubation of new companies, larger public-sector projects, and/or creation of facilities designed to host and maintain FRO outputs (e.g., datasets or tools).

Mission selection

FROs should pursue specific goals that, if achieved, will dramatically increase the R&D capacity and/or technological capabilities of the United States in a given field. To preserve the FRO program's ability to pursue specific, focused innovation objectives, FROs would operate for defined time periods and would not ordinarily be renewed. Renewal would only be permitted in exceptional cases in which an FRO proves that an extension of that FRO would be as impactful as the initial investment. More frequently, we expect that an FRO might serve as proof of concept for a project or initiative that could then be separately pursued through an act of Congress or through a public-private partnership. All new FROs should meet following two criteria:

(1) **FROs should be transformative.** While FROs might occasionally integrate existing methods to directly produce a new dataset or clinical/scientific outcome, FROs should generally focus on developing transformative new technologies, systems, or processes. These capabilities should reduce the cost and/or increase the speed and reliability of subsequent scientific, clinical, or other downstream efforts, substantially increasing the rate of overall science and technology development in the United States.

(2) **FROs should be focused.** Each FRO should be established with a clear, goal-oriented purpose. FROs should driven by quantitative metrics and/or concrete design goals and should be limited in scope and duration. Serendipitous discoveries made during the course of FRO research that are outside of the mission scope should be shared freely with external researchers for follow-up. Though we expect FROs to work closely with universities, FROs must not become subject to academic incentives and must avoid mission creep. Although an FRO may maintain external (e.g., academic) advisors and consultants, core staff must be appointed full-time at the FRO.

To ensure efficient and decisive selection and oversight of FROs, a dedicated and innovative program manager—rather than a committee of peer reviewers—could be recruited to help drive

the conception, selection and formation of a small number of FROs on the government side. DARPA similarly appoints program managers instead of committees to enable the embrace of visionary or divergent perspectives. Program managers should be willing to take risks on "moonshot" projects for which there is not a consensus on feasibility or likely value.

Frequently Asked Questions

What new value would FROs bring to our existing R&D establishment?

The table below shows how FROs fill a key organizational gap and complement existing R&D structures in the United States.

Characteristic	FRO	DARPA (or other ARPAs)	National Lab	Academia	Startup	Corporate Lab
Generates public goods	Yes	Yes	Yes	Yes	No	Limited by shareholder incentives
Operates in the pre-commercial phase	Yes	Yes	Yes	Yes	No	Limited by shareholder incentives
Conducts large-scale team science, where no single researcher receives primary credit	Yes	Limited with academic performers, easier with corporate performers	Yes	No	Yes	Yes
Is managed/led by a CEO, rather than existing as a collaboration	Yes	Finite ability of program manager to fully align incentives of performers	Yes	No	Yes	Yes
Requires scalable, repeatable execution, not simply "proof of concept"	Yes	Limited with academic performers, easier with corporate performers	Yes	No	Yes	Yes
Requires primarily professional staff who may need industry-level compensation, not just graduate students and postdocs	Yes	Limited with academic performers, achieved with corporate performers	Yes	No	Yes	Yes
Involves a high level of accountability for achieving defined, time-bound milestones	Yes	Limited with academic performers and with certain industry performers	Yes	No	Yes	Yes
Focuses on conceptual breakthroughs and/or open-ended exploration	No	Yes	Not typically	Yes	No	Depends on company and lab
Emphasizes collaborations among widely distributed groups with divergent incentives and cultures	No	Yes	No	No	No	No
Exists as an autonomous organization mobilized in a rapid, agile fashion	Yes	Yes (inasmuch as distributed programs constitute autonomous organizations)	No	No	Yes	No

Focused on training students	No	Only to the degree that specific programs are executed in part by academics	No	Yes	No	No
Is a permanent institution/career path	No	No	Yes	Yes	No	Yes
Provides strong support for post-project transition to commercialization	Yes	In some cases	No	No	Yes	In some cases

What are concrete examples of FROs?

The Human Genome Project is a good example of a past effort that could have easily existed as an FRO. A Genome Sequencing FRO could have been established to develop key genome sequencing instrumentation, processes, and protocols. The success of the FRO could have been assessed using metrics such as "letters of DNA sequenced per day" with an appropriate means of quality cross-checking. Once that metric exceeded a pre-defined threshold, the FRO would have been deemed to have achieved its goal. The sequencing capacities that the FRO developed would have then been used by other institutions and facilities to contribute to the much larger Human Genome Project itself. The appendix provides additional concrete examples of areas that could benefit from FROs.

Why can't the projects envisioned for FROs be pursued in academia or industry?

Academics are rewarded for novelty and individual publication records, which discourages the tight-knit collaboration (among potentially dozens of people) and accountability necessary for large-scale projects to succeed. For-profit companies can align large and well-organized teams but must pursue developments that are commercially profitable and achievable within the lifetime of a patent and/or the timescales dictated by investor returns. Raising the amount of money needed for FRO-level technical work in industry would, except in the rarest of cases, require an immediate path to revenue or profitability. Many societally important scientific and technological projects take too long for venture-capital funding cycles, and/or create public goods without the ability to efficiently capture value from them. To highlight these points further:

FRO projects are generally not doable in academia because their requirement for coordinated, milestone-driven team-based engineering is not consistent with academic culture and reward structure. Instead, the work done in FROs requires:

- Large-scale team science, rather than a study led by a single principal investigator or a small group of co-researchers.
- Top-down coordination and management/leadership from a CEO/CTO figure, rather than a loose collaboration between peers with limited direction or long-term commitment.
- Scalable, repeatable execution, rather than simply "proof of concept" of a new idea.

- Involvement of highly trained and experienced professionals, rather than primarily graduate students and postdocs.
- A high level of accountability for achieving time-bound milestones.
- Goal-driven development, integration and systematization of technologies, rather than prioritization of unpredictable conceptual breakthroughs and/or open-ended exploration.

FRO projects are not generally doable in industry because they have no immediate (within about five years) path to profitability. Instead, the work done in FROs requires:

- Generating scientific data that will be made freely available to the research community.
- Creating a new scientific instrument that will be transformative for research in the longer term but may not generate direct financial returns over relatively short timeframes.
- Developing crucial technology that is still immature and may require sustained investment before becoming commercially viable.

Why is roughly $5–15 million per year per FRO an appropriate amount of funding?

This amount is roughly comparable to the funding level of a Series A or B level hard-tech startup or a DARPA project. $5–15 million per year over five years will support a team of a dozen or more highly skilled senior staff (recruited with industry-level salaries) as well as requisite capital infrastructure, logistical support, and workspace. A robust funding commitment is also needed to provide a compelling reason for top talent to leave their regular jobs and to finance a large-scale project that is infeasible in an academic or for-profit setting. A high funding level is justified by the capabilities that FROs will develop. The tools, technologies, datasets, and processes developed by FROs will substantially reduce the cost or increase the throughput or reliability of subsequent work, enabling much larger returns on investment for the long term.

Why can't existing funding structures create centralized, cross-disciplinary projects in an agile way?

Existing funding structures, such as grant programs administered by federal agencies, tend to distribute work across many different institutions like universities. These institutions are generally further fragmented into many small labs—each of which is pursuing a different problem and training their own students under an individual principal investigator. This system is useful for studying a variety of specific components of a given field but fails at efficiently tackling many large, complex technological or measurement challenges. While such challenges may be best addressed through a single, coordinated $50 million-dollar scale project with an optimal team structure, existing funding structures will allocate that $50 million across a hundred projects of $500,000 each. The result is often a focus on components rather than systems, under-powered studies, partially duplicated yet incompatible efforts, and fragmented incentives among transient or inexperienced staff. Even ARPA organizations rely on external entities as performers and can suffer when no single organization is both able and willing to pull the pieces together. Note that

there are notable exceptions, such as when DARPA convinced the company Moderna to pivot some of its efforts from cancer to vaccine development, and the ARPA model has complementary advantages to FROs for certain problem classes.[2]

When teams are not aligned around a single mission and led by a CEO figure (possessing both technical and management experience), there is limited accountability, transparency, and collaboration. This truth is well-known in business but not generally acknowledged in government-funded R&D. The separation of agencies like NIH into institutes or divisions focused on individual diseases also contributes to difficulty in assembling centralized teams around cross-cutting initiatives. DOE National Labs are a rare example of a federal institution that does succeed at pursuing centralized large-scale R&D projects. However, DOE National Labs naturally focus on the energy sector and are mostly characterized by long-term contracts focused on fixed infrastructure assets, rather than agile development of new methods.

How are FROs different from DARPA?

Multiple aspects of DARPA provide inspiration for FROs, including the milestone-driven and ambitious, risk-taking nature of DARPA programs, DARPA's ability to recruit top talent for intensive projects of finite duration, and DARPA's intermediate scale that bridges a gap between typical grants and very large initiatives. Indeed, some FROs could be created within or as an outcome of broader DARPA or DARPA-like programs. Certain FROs could also coordinate DARPA-like externalized research if needed and appropriate. Yet FROs would depart from DARPA in two important ways:

1. DARPA focuses on defense-specific applications. FROs would focus on producing public goods for science, technology, or medicine broadly.
2. DARPA relies on portfolios of externalized research through transiently organized collaborations of existing organizations. FROs would be dedicated, stand-alone organizations led by a CEO-like figure and staffed with the personnel needed to carry out projects in house.

How are FROs different from National Lab projects?

DOE National Labs are world-class at building large-scale user facilities and at maintaining high-security technical infrastructure. FROs would largely pursue projects that don't require these specific capacities and/or are out of scope topically for the DOE labs. FROs would also have more open and fluid interactions with industry and academic labs compared to DOE labs and could be operated at lower security levels.

[2] Washington Post, Paul Sonne, "How a Secretive Pentagon Agency Seeded the Ground for a Rapid Coronavirus Cure," (September 22, 2020). https://www.washingtonpost.com/national-security/how-a-secretive-pentagon-agency-seeded-the-ground-for-a-rapid-coronavirus-cure/2020/07/30/ad1853c4-c778-11ea-a9d3-74640f25b953_story.html.

About the Authors

Sam Rodriques is an entrepreneur, biotechnologist, and inventor. He is founding the Applied Biotechnology Laboratory at the Francis Crick Institute in the winter of 2021. He has founded multiple medical-device companies and obtained his Ph.D. in Physics from the Massachusetts Institute of Technology as a Hertz Foundation Fellow.

Adam Marblestone is a Schmidt Futures Innovation Fellow and a Senior Fellow with the Federation of American Scientists. He was previously a research scientist at Google DeepMind, was the Chief Strategy Officer of the brain-computer interface company Kernel, and co-founded BioBright LLC. He has been recognized with a Technology Review 35 Innovators Under 35 Award. He received a Ph.D. in Biophysics from Harvard University as a Hertz Foundation Fellow.

About the Day One Project

The Day One Project is dedicated to democratizing the policymaking process by working with new and expert voices across the science and technology community, helping to develop actionable policies that can improve the lives of all Americans, and readying them for Day One of a future presidential term. For more about the Day One Project, visit dayoneproject.org.

Appendix: Conceptual Ideas for FROs

High-throughput brain mapping

Many neuroscientists believe that all mammals, including humans, share fundamental features of a core cortical architecture that underlies cognition. Knowing that architecture is crucial for understanding brain disease and for generating the next generations of artificial intelligence and energy-efficient computing. In an effort largely executed by the Allen Institute for Brain Science, the Intelligence Advanced Research Projects Activity (IARPA) MICRONS program brought brain-circuit mapping to the 1 cubic millimeter (mm^3) level. Yet the brain of a mouse is about 400 times larger than the volume mapped through the MICRONS program. The NIH is considering a transformative project to map the mouse brain, but the requisite tools have not yet been brought together to support this effort.

An FRO focused on brain mapping could advance and validate the technologies necessary for mapping an entire mammalian brain. A specific governing metric could be to "achieve a 5x increase in mapping throughput at a 5x reduction in cost, relative to current methods." Talent could be drawn from the industrial process control and semiconductor inspection industries, from tech companies, and from academic neuroscience and chemistry departments. Resulting systems could then be internalized into a National Lab and scaled further.

Multi-drug combination therapies and drug repurposing

Neither academics nor pharmaceutical companies are well incentivized to pursue systematic studies of multi-drug combination therapies, especially those that would use combinations of already-approved drugs to optimize long-term health outcomes in a preventative fashion. Academic labs are rewarded for novelty and specialization and are under pressure to build a unique academic "brand" around a particular mechanistic biology "story" rather than for applied engineering or logistical achievements. Pharmaceutical companies need to get single, novel, and patentable drug assets through trials with minimal risk. Yet combinatorial interventions could do much to crack the complex and interlocking problems of age-related disease, health span, and cancer.

An FRO could centralize resources around combinatorial testing of compounds already known to be safe. The FRO could specifically be tasked with developing combinatorial interventions that prevent the onset of diabetes or heart disease in at-risk populations. If successful, the FRO's combinatorial approach would likely be applied more broadly to address other pressing health issues. Talent could be drawn from the pharmaceutical industry, medical centers, and from engineers with expertise in robotics and real-time controls.

Synthetic fuel production

Many researchers are working on new methods to produce synthetic, carbon-neutral fuels. This requires integration of innovations in catalysis and materials with state-of-the-art engineering methods for energy storage and fuel transport. Unfortunately, today the synthetic-fuels field is

fragmented among many academic researchers pursuing small-scale studies outside of relevant systems (e.g., fuel markets, logistics operations). An FRO could seek to integratively solve the problem of synthetic fuel production and storage by bringing top chemists and materials scientists into a system-building context that emphasizes integrated devices rather than just component materials and is equipped with high-throughput robotic materials screening and computational modeling capabilities, not just chemistry labs. Talent could be from fuel-transport and logistics companies, battery manufacturers, and academic experts in materials science, electrical engineering, and chemistry.

Diagnostics and biomarkers for priority diseases

It is hard to develop therapeutics for age-related diseases like Alzheimer's in the absence of good early biomarkers. Developing biomarkers and early diagnostics can be prohibitively expensive for academic researchers since, due to the low incidence and slow progression of any one disease in the general population, it requires very large and well-controlled studies to prove statistical significance. Better-resourced pharmaceutical companies are, meanwhile, not well incentivized to develop biomarkers in the absence of a marketable treatment. A similar market failure exists for early cancer detection.

An FRO could work to collect a large pool of clinical samples to support development of new biomarkers for a disease like Alzheimers. The FRO could also use machine learning on large anonymized patient populations to validate the resulting biomarkers at a level of statistical power, control, and reproducibility inaccessible to standard academic labs. These biomarkers could then be made publicly accessible via commercialization. The underlying systems for developing new biomarkers could be separately commercialized and/or internalized as an NIH intramural center.

Developing next-generation microprocessors

Microprocessor advancement is limited by the number of transistors that engineers can pack onto microchips of a fixed size. The next generation of microprocessors will need to rely on alternative strategies, such as using nanophotonics to quickly move data between separate chips and/or to package chips tightly in 3D.

An FRO could be tasked with developing a prototype large-scale 3D photonic interconnect system to support next-generation microprocessors. The FRO would bring top optical physicists from academia together with chip designers and machine learning experts from industry. The FRO would contract with commercial manufacturers for prototype fabrication at a level inaccessible to academia and too capital-intensive for a seed stage startup, while leveraging the latest research ideas from multiple fields to inform prototype architecture. If successful, a large semiconductor or tech company might invest in subsequent commercialization and optimization of the prototype. A DARPA or NSF program split across several university labs and small startups could support development of software algorithms compatible with the new hardware to further pave the way for dramatic advances in computing.

Organ preservation and banking

Solving the organ transplant shortage would save more lives each year than curing cancer.[3] In theory, rapidly cooling and reheating organs or biochemically inducing a state of suspended animation should allow organs to be maintained for much longer times to facilitate organ banking and transport. This problem is currently too high-risk and multifaceted for startups to tackle, especially given regulatory risks. Making organ preservation and banking a reality therefore requires a large centralized multidisciplinary engineering effort closely linked to clinical settings.

An FRO could locate an engineering and operations center adjacent to a major hospital system and recruit top cryogenics and microwave physics experts as well as experts in epigenetic tissue characterization and bioreactors to develop an organ preservation and storage pipeline that enables organs to last at least 10 times longer than is currently possible. NIH could later initiate a program to fund clinical trials of this system for multiple organ types.

[3] Giwa et al., "The promise of organ and tissue preservation to transform medicine," (2017) https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5724041/.