

Cancer: A Computational Disease That AI Can Cure

Jay M. Tenenbaum and Jeff Shrager

■ *Cancer kills millions of people each year. From an AI perspective, finding effective treatments for cancer is a high-dimensional search problem characterized by many molecularly distinct cancer subtypes, many potential targets and drug combinations, and a dearth of high-quality data to connect molecular subtypes and treatments to responses. The broadening availability of molecular diagnostics and electronic medical records presents both opportunities and challenges to apply AI techniques to personalize and improve cancer treatment. We discuss these in the context of Cancer Commons, a “rapid learning” community where patients, physicians, and researchers collect and analyze the molecular and clinical data from every cancer patient and use these results to individualize therapies. Research opportunities include adaptively planning and executing individual treatment experiments across the whole patient population, inferring the causal mechanisms of tumors, predicting drug response in individuals, and generalizing these findings to new cases. The goal is to treat each patient in accord with the best available knowledge and to continually update that knowledge to benefit subsequent patients. Achieving this goal is a worthy grand challenge for AI.*

Despite enormous strides in scientific understanding and medical treatment, cancer still kills millions of people each year. This article will explore why cancer has proven to be such a formidable adversary, and how AI and machine learning (ML) can save lives by helping individual patients beat the odds.

Modern molecular biology supports the hypothesis that cancer is actually hundreds or thousands of rare diseases, and that every patient’s tumor is, to some extent, unique (Pleasant et al. 2010). Although there is a rapidly growing arsenal of targeted cancer therapies that can be highly effective in specific subpopulations,¹ especially when used in rational combinations to block complementary pathways, the pharmaceutical industry continues to rely on large-scale randomized clinical trials that test drugs individually in heterogeneous populations. Such trials are an extremely inefficient strategy for searching the combinational treatment space, and capture only a small portion of the data needed to predict individual treatment responses. On the other hand, an estimated 70 percent of all cancer drugs are used off label in cocktails based on each individual physician’s experience, as if the nation’s 30,000 oncologists are engaged in a gigantic uncontrolled and unobserved experiment, involving hundreds of thousands of patients suffering from an undetermined number of diseases.² These informal experiments could provide the basis for what amounts to a giant adaptive search for better treatments, if only the genomic and outcomes data could be captured and analyzed, and the findings integrated and disseminated.

Toward this end we are developing Cancer Commons,³ a family of web-based rapid-learning communities in which physicians, patients, and scientists collaborate to individualize cancer therapy (Shrager, Tenenbaum, and Tavers 2011). The goals of this initiative are to (1) give each patient the best possible outcome by individualizing his or her treatment based on their tumor’s genomic subtype, (2) learn as much as possible from each patient’s response, and (3) rapidly disseminate what is learned. The key innovation is to run this adaptive search strategy in “real time,” so that the knowledge learned from one patient is disseminated in time to help the next.

Cancer Commons provides an exciting and important

domain for the application of AI and ML. Achieving its goals will require developing models for predicting individual responses to cancer therapies and using them to plan and run thousands of adaptive treatment experiments. The genomic and clinical response data from each patient can be analyzed to continuously improve the predictive models, as well as our understanding of cancer biology and drugs. Each patient is treated in accord with the best available knowledge, and that knowledge is continually updated to benefit the next patient.

In the next section, we review the current paradigm for translational cancer research, and its limitations, as we enter the new era of genomics and personalized medicine.⁴ Then we introduce Cancer Commons as a new, open-science paradigm for real-time translational research, with no daylight between lab and clinic. The fourth section explores the opportunities and challenges at the intersection of Cancer and AI research. Finally, we conclude with an invitation to participate in the grand challenge of curing cancer.

Cancer Research in the Era of Genomics and Personalized Medicine

It is unlikely that a single cure will be found for cancer—it is not a single disease. To a first approximation, cancer is what happens when a normal cell mutates in such a way that it can reproduce but no longer control its growth. As genome sequencing becomes less expensive and more widespread, scientists are finding a spectacular variety of mutations that lead to cancerous growth (Pleasant et al. 2010). Consensus is building that, because cancer is an individualized disease, it must be addressed through individualized treatment (see, for example, McDermott, Downing, and Stratton [2011]).

Developing and justifying individualized cancer treatments is a very difficult problem. First, each organ-specific cancer involves many rare genomic alterations, thus leading to highly heterogeneous disease populations. This heterogeneity flows directly from the underlying biological and genetic complexity of cancer. The mutation history of each tumor cell is also unique—the result of cascading and compounding errors in the cell's machinery for maintaining the integrity of its DNA, as well as evolutionary adaptations that help it evade therapy.

Genetic mutations typically manifest themselves through cell signaling pathways, chains of protein-catalyzed reactions that govern basic cellular functions such as cell division, cell movement, how a cell responds to specific external stim-

uli, and even cell death. The proteins in a pathway often communicate by adding phosphate groups, which act as “on” or “off” switches, to a neighboring protein. When one of the oncogenes that codes for such proteins is mutated, the resulting protein can be stuck in the “on” or “off” position, leading to uncontrolled cell growth, and thus cancer.

In contrast to traditional chemotherapy, which targets all rapid cell growth, modern targeted cancer therapies attempt to block specific molecules in the signaling pathways. Targeted therapies may be more effective than current treatments and less harmful to normal cells. However, a given patient may require a custom cocktail of targeted therapies to block both the pathways directly affected by an oncogene as well as secondary pathways that may enable the cell to evade therapy.

Large-scale clinical trials are problematic in genomic diseases like cancer because they rely on population statistics rather than individual responses. A drug that works on 50 percent of patients tested may or may not be better for any given patient than one that works on 20 percent. This inability to account for individual responses may explain why so many late-stage trials fail to demonstrate statistical efficacy, even though a few individuals do respond. What disease did these responders have, that used to be called “breast cancer” or “melanoma,” for which there is now an effective drug? How many other cancer patients have the same genomic disease? Unfortunately, we'll never know because when a trial fails, the drug is, as a rule, abandoned.

One happy exception to this rule involves Gleevec (Imatinib), originally developed in 1998 by Novartis for chronic myelogenous leukemia (CML). Gleevec was among the first of a new generation of targeted therapies developed to block a specific enzyme—in this case, BCR-ABL—in a cancer signaling pathway. Shortly after its development, it was experimentally observed that Gleevec also blocked a structurally similar enzyme called c-KIT, which was a primary cause of gastrointestinal stromal tumors (GISTs), a rare form of sarcoma (Demetri 2002). Gleevec was subsequently tested on many forms of cancer, including a 2003 trial for melanoma that failed (Kim et al. 2008). There were, however, a handful of responders in that trial, all of whom had two things in common: their primary tumors were located not on their skin, but in places where the sun don't shine (for example inside the mouth), and their tumors all had a mutated c-KIT gene. In 2008, results of a definitive trial testing Gleevec on patients whose tumors had been pretested for c-KIT mutations demonstrated a major response (Hodi et al. 2008). A randomized trial would have needed perhaps a million patients to demonstrate statistical significance for a drug that affected at most 3 percent of melanoma

patients. There's more to this story, and we'll return to it several times.

As we enter the age of personalized genomic medicine, the Gleevec experience will become the norm. At the molecular level, cancer is hundreds if not thousands of unique diseases. Given that there are more than 800 therapies in development targeting specific mutations, and that they will often have to be used together in cocktails to ensure a durable response, there simply isn't enough time, money, patients, or specimens to test treatments individually on large heterogeneous patient populations. We need a smarter way to search for effective treatments, one that utilizes the complete molecular and clinical profile of every patient to efficiently decide which drugs are likely to work best in a specific patient.

Community oncology practices, where the vast majority of patients are treated, provide a promising way forward. Every day, thousands of patients who have exhausted the standard of care are treated with off-label drugs and cocktails. These treatment decisions are based largely on the judgments and experience of individual physicians, and the results are seldom reported. What if we could coordinate these thousands of ad hoc "N of 1" experiments, capture their results, and integrate them with the evidence from large-scale controlled trials? Cancer Commons provides the platform for doing just that, transforming the everyday practice of oncology into a giant adaptive search for better treatments.

Cancer Commons

Cancer Commons is a family of web-based rapid learning communities where physicians, scientists, and patients collaborate to individualize therapies based on the molecular profile of each tumor. Cancer Commons provides physicians with the knowledge and resources to treat each patient in the best possible way, and to disseminate the resulting data and knowledge as efficiently as possible to help subsequent patients. Cancer Commons is, in essence, a huge adaptive clinical trial, whose goal is to continuously improve patient outcomes, rather than

to demonstrate the efficacy of a specific drug for regulatory approval.

Cancer Commons is being developed one cancer at a time, beginning with melanoma. At the core of each Commons is a molecular disease model (MDM). MDMs are expert-curated reference models enumerating the known molecular subtypes of a cancer. Each subtype is linked to relevant pathways, diagnostic tests, approved and experimental (targeted) therapies, and clinical trials. At each point in time, oncologists can use the appropriate MDM to treat patients with the best available therapies for their tumors' molecular subtype. The subtypes and associated therapies can be continually refined based on how patients respond. For example, subtypes can be split corresponding to responders and nonresponders, and new subtypes added to accommodate previously unseen tumor types. Conceptually, all patients within Cancer Commons are participating in a huge, continuously running, adaptive clinical trial that is constantly testing and refining both the MDM subtypes and their potential treatments.

The power in the Cancer Commons approach lies in the rapid feedback loop that is generated when a patient is rationally treated based on his or her molecular subtype but does not respond as predicted. Armed with state of the art genomics technology, researchers can attempt to uncover why that individual failed to respond and apply the findings to the next patient with the same molecular subtype. Studying response outliers can be particularly informative. Earlier, we discussed the example of a handful of melanoma patients who responded to Gleevec in a trial that failed to meet its clinical end points. The flip side is the subset of patients who fail to respond in a successful trial. For those patients that were directed to the trial through Cancer Commons, one can go back to the scientists on whose research the molecular disease model was based and ask them to revise the model based on the new molecular and clinical data from the subset of nonresponding patients. What disease subtype did they have and how should it be treated? These questions are not academic;

they pertain to patients who urgently need new options. Over time, the answers should lead to MDM models with increasingly specific molecular subclasses, linked to increasingly efficacious therapies.

Cancer Commons' rapid learning loop is illustrated in figure 1. It applies the same continuous improvement philosophy to drug development and clinical trials that management gurus like W. Edwards Deming used in the 1980s and 1990s to slash time and costs in automobile manufacturing and many other industries. It is also related to the concept that Andy Grove, the former CEO of Intel, had in mind when he advocated, in an influential *JAMA* editorial, that the pharmaceutical industry should strive to achieve the fast "knowledge turns" achieved by chipmakers (Grove 2005). According to *The Economist*, "A knowledge turn [...] is the time it takes for an experiment to proceed from hypothesis to results, and then to a new hypothesis—around 18 months in chipmaking, but 10–20 years in medicine."⁵

In addition to accelerating clinical research, Cancer Commons also accelerates communication of the results and subsequent collaborations. The MDMs serve as living review articles, maintained online and continuously updated by top cancer experts. Clinicians and researchers can post peer-reviewed clinical observations and data, anecdotal case reports, rational treatment hypotheses, and actionable research findings to the relevant MDM subtype—preliminary results that are too early for formal presentation in journals or at conferences but that may be helpful to late-stage patients running out of options. Physicians and patients who choose to act on these postings can then report clinical outcomes and side effects from their "N of 1" experiments through the web. Treatment hypotheses that prove successful in a few patients can be generalized to a specific molecular phenotype and validated in proof-of-concept studies or small clinical trials. Hypotheses that hold up can ultimately be incorporated into the guidelines as the new standard of care. Hypotheses and observations that don't hold up can be quickly rejected or revised based on the data points from human subjects.

A Melanoma Molecular Disease Model (MMDM)

Figure 2 depicts a portion of the current MDM for melanoma, created by CollabRx, Inc. (Vidwans et al. 2011) This expert-curated model summarizes what is currently known about each molecular subtype of melanoma—which biochemical pathways are disrupted, their prevalence, their genetic markers, and, critically, how they respond to therapeutic interventions.⁶ Figure 2a illustrates the major signaling pathways implicated in melanoma, including the map kinase (MAPK) pathway (on the left) and the AKT/PI3K pathway (on the right), which regulate cell growth, cell proliferation, and cell death. There is a lot of crosstalk between these pathways and their downstream effectors. Figure 2b is a table summarizing the subtypes of melanoma and linking them to these pathways. Each subtype is defined based on the mutational status of a key oncogene or tumor suppressor gene in a given pathway (such as BRAF for subtypes 1.1 to 1.4 and c-KIT for subtype 2.1). Each subtype is then further divided by mutations that may coexist in genes that play a supportive role (such as *PTEN*, *AKT*, and *CDK4* in the case of subtypes 1.2, 1.3, and 1.4). The subtypes are ordered by their prevalence and potential for therapeutic intervention. All told, the model currently distinguishes eight actionable subtypes. Melanomas that do not fit into these subtypes are parked in subtype 9 (TBD) pending further research.

The MMDM represents the current state of knowledge about melanoma subtypes and their associated molecular diagnostics and targeted therapies. It connects information about pathways, targets, tests, drugs, and trials previously scattered across hundreds of publications and databases, not to mention information that previously existed only in the heads of a few experts. It is clearly a work in progress. As oncologists use the MDM to treat patients, the subtypes and their associated therapies will be continually refined based on how individual patients respond. Subtypes will be split corresponding to responders and nonresponders, and new subtypes added to accommodate previously unseen tumor types. Over time, it would not be surprising to see the model grow to encompass dozens of subtypes and sub-subtypes with distinct therapeutic options.

The MMDM appeared in a peer-reviewed scientific journal (Vidwans et al. 2011). The PDF, containing many tables, charts, literature references, and supporting data, is also available online,⁷ hosted on a semantic media wiki (SMW) platform.⁸ SMW is an extension of MediaWiki—the wiki application best known for powering Wikipedia. The semantic extensions transform the MMDM into a hyperdocument that can be rapidly navigated to find all information relevant to a given topic

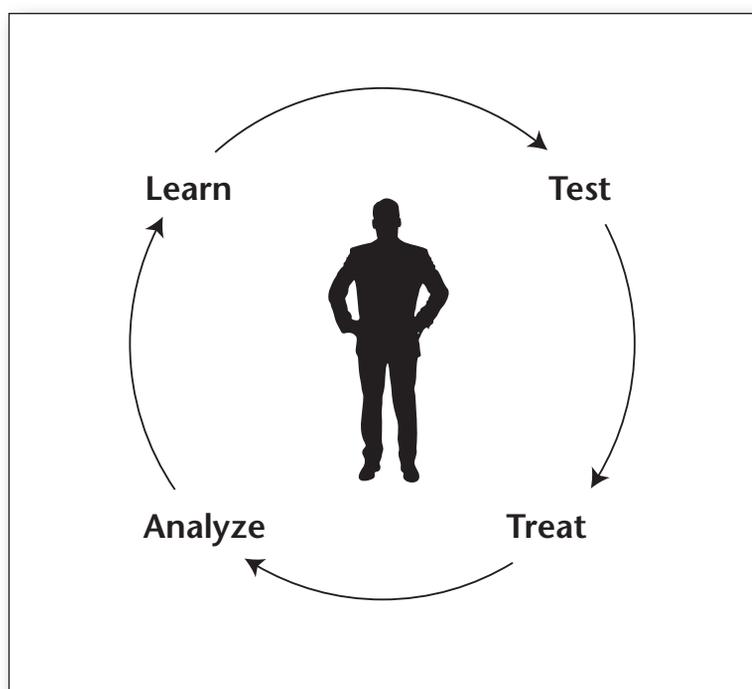


Figure 1. Cancer Commons “Rapid Learning” Loop.

(such as subtype, test, drug, or trial) as well as additional information on the web. Moreover, the SMW adds semantic annotations that make the content of the MMDM understandable by a computer. The MMDM can thus function as a collaborative database, whose contents can be accessed through the semantic web.

While valuable to medical and AI researchers, few practicing physicians or patients are likely to use the MMDM in either of the above forms to guide treatment decisions. What is needed are user-friendly web-based applications that transform the knowledge in the MMDM into personalized, actionable information.

The Cancer Commons Platform and Applications

Cancer Commons participants can interact with the molecular disease model through web-based and mobile applications (apps), each designed to do one task well (see figure 3). The first such app is a targeted therapy finder (TTF) for melanoma. It was developed by CollabRx, Inc.,⁹ and released at the end of 2010 (see figure 4). The CollabRx TTF guides advanced melanoma patients and their physicians through a series of molecular tests to determine their tumor’s subtype, and then describes experimental therapies based on the best available knowledge. Additional apps, under development, will support outcomes capture and analysis, adaptive treatment planning, and collaborative updating of the MMDM based on the latest clinical and laboratory findings.

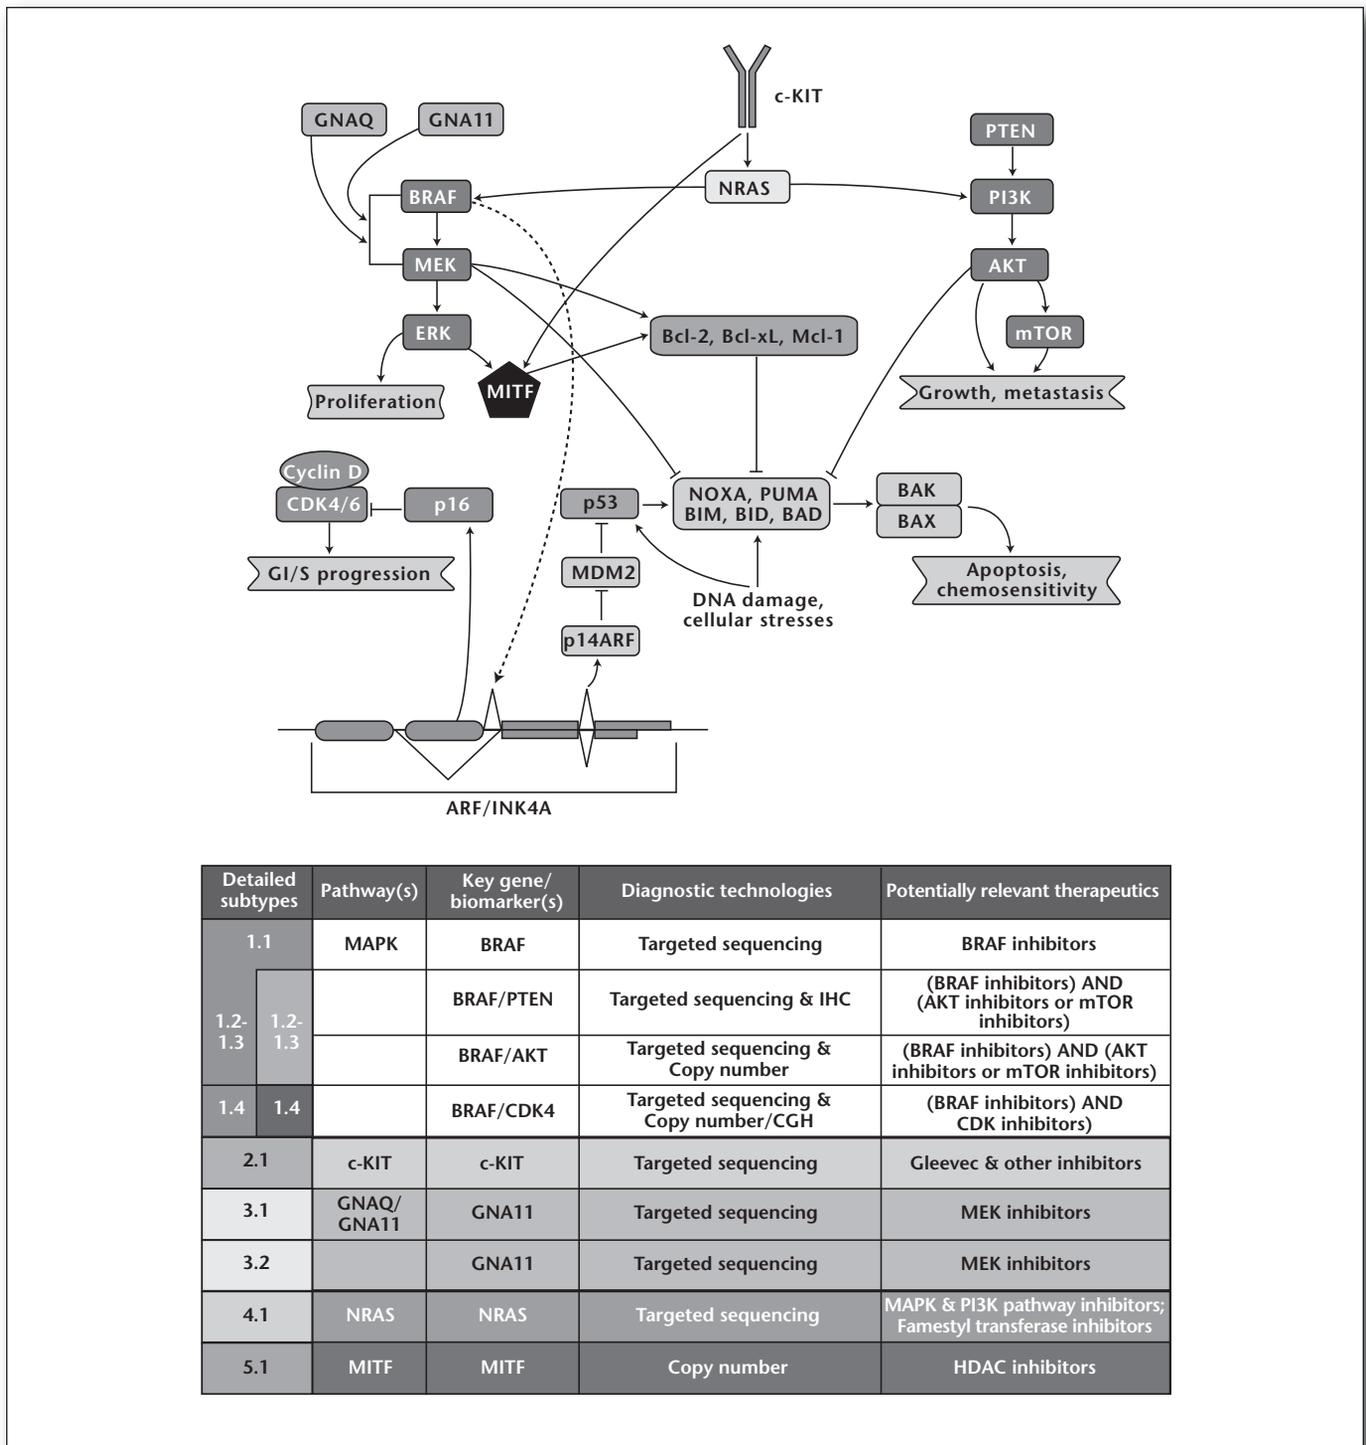


Figure 2. A Portion of the Current MDM for Melanoma, Created by CollabRx, Inc.

(a) A highly simplified depiction of cell signaling pathways known to be active in melanoma. The MAPK pathway communicates a signal from growth factor receptors on the surface of the cell to the DNA in the nucleus of the cell, triggering cell division. A defect in this pathway can lead to uncontrolled cellular proliferation. The PI3K pathway is involved in a range of cellular functions implicated in cancer, including cell growth, proliferation, differentiation, motility, and survival. Interactions between the MAPK and PI3K/AKT pathways lead to proliferation and survival of melanoma cells. The CDK pathway, shown in blue, regulates progression through the cell cycle, which includes stages for DNA replication, inspection (checkpoints), and repair. The P53 pathway plays several critical anticancer roles, initiating DNA repair when DNA has sustained damage, and apoptosis (programmed cell death) if the damage proves to be irreparable. One or more of these pathways are mutated in the vast majority of melanoma cells. (b) A subset of the table relating the biochemical pathways to subtypes of melanoma and their possible treatments.

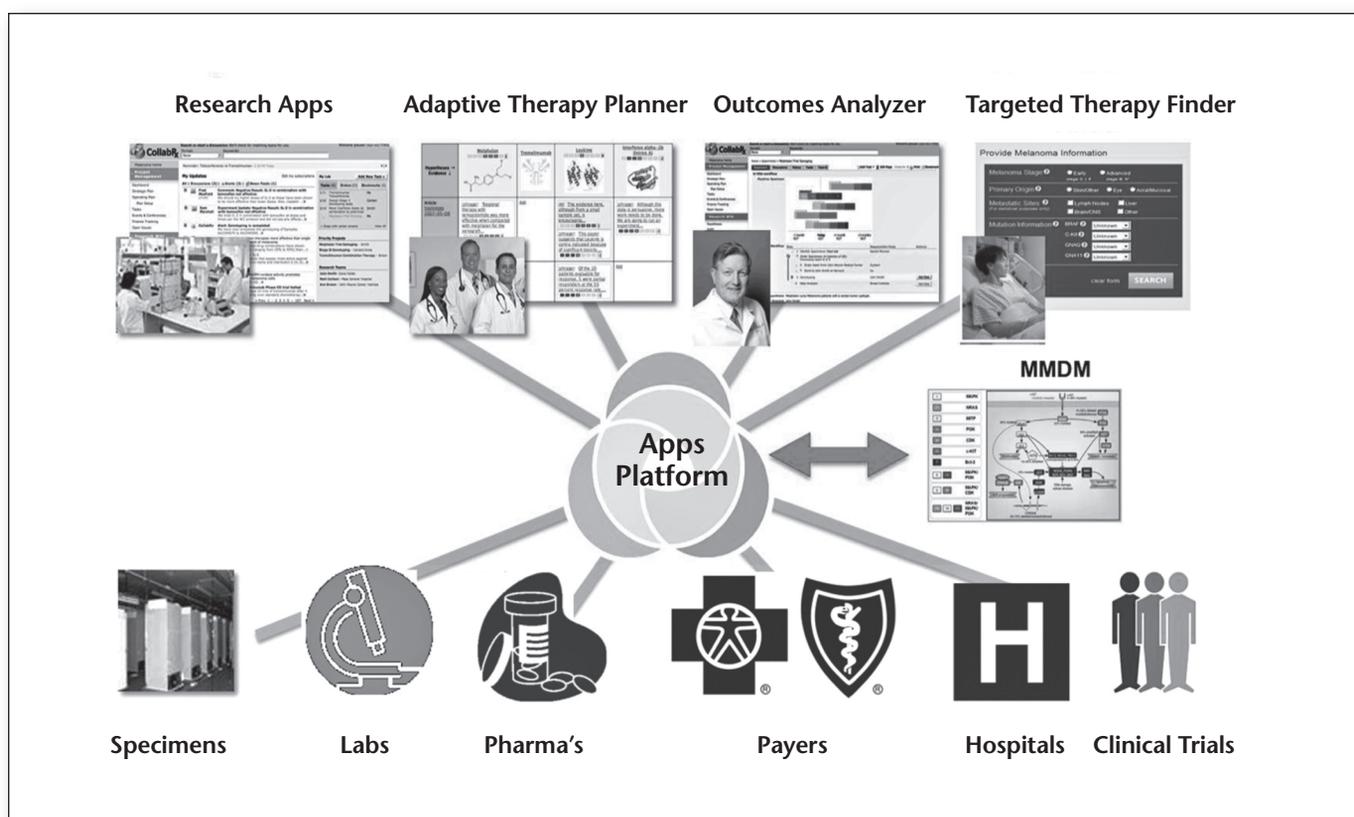


Figure 3. Schematic of the Cancer Commons Platform.

At the center is an applications platform and knowledge hub that draws from, and feeds, the semantically represented MDM (center right). Specific applications (and example users) are depicted at the top, and examples of services are depicted along the bottom.

The Cancer Commons application platform has open APIs for third-party developers to create sophisticated medical decision support apps that, like the TTF, combine the science in the MMDM with relevant clinical and personal considerations (for example, the patient's state of health and response to prior treatments, the availability of reimbursement). Since the platform is web based, such apps can also access a rich ecosystem of medical and research resources, including molecular diagnostics laboratories, specimen and data repositories, and high-throughput drug screening and animal testing facilities.

The Cancer Commons platform—encompassing the MDMs, apps, and services ecosystem—provides the computational framework for the rapid learning loop illustrated in figure 1. Expert hypotheses about subtypes and their potential treatments, codified in the MMDM, can be tested in the clinic as appropriate patients present. Hypotheses can also be tested in the laboratory on appropriate cell lines or animal models. The resulting clinical and laboratory observations may give rise to new hypotheses, which can be vetted by the experts and added to the model for future validation.

The Cancer Commons Network

Cancer Commons is being developed one cancer at a time in partnership with leading professional and patient advocacy organizations, pharmaceutical companies, medical centers, and health informatics companies. These disease commons form a network (figure 5) that will exploit the significant economies of scale and opportunities for cross learning (for example, common pathways and drugs) that come from rethinking cancer as a molecular disease.

Following the process developed in melanoma, each new Commons begins by recruiting a panel of experts. The panel develops an MDM for that cancer. The knowledge in the MDM is then represented in a semantic media wiki, so it can be accessed through apps. Several new commons are currently in various stages of development, including sarcoma, breast, and lung cancer.

Linking these Commons into a knowledge network occurs naturally through shared concepts such as subtypes, pathways, targets, tests, drugs, and trials. For example, while many cancer pathways have been identified, a dozen or so appear over and over again, including the MAPK, PI3K,

Figure 4. Screen Shot of the Melanoma Targeted Therapy Finder App¹⁰
Developed by CollabRx, Inc., Based on the Melanoma MDM.

Selecting characteristics of a patient's specific tumor and pressing "SEARCH" will report clinical trails and literature that are possibly appropriate for this individual patient.

CDK, c-KIT, and P53 pathways. These play a prominent role in melanoma and many other solid tumors. When cancers share a common pathway or subtype, much of the information, experimentation, and hard-won research results can potentially be leveraged across them. In particular, targeted therapies that prove successful in one of these cancers might also work in the other. The aforementioned example of Gleevec is an excellent case in point; its effectiveness against c-KIT mutations was first demonstrated on GISTs (a form of sarcoma), and it was subsequently successfully applied to c-KIT mutated melanomas. The hope and expectation is that many such synergies will emerge from Cancer Commons' knowledge network—slashing years off the decade-long process it took to establish the efficacy of Gleevec—first for GISTs, and then for melanoma.

Many experts believe that cancers will one day be characterized as molecular diseases (for example, c-KIT or BRAF disease) rather than organ-based ones. The Cancer Commons network could accelerate that future.

AI Opportunities and Challenges

From an AI perspective, finding effective treatments for cancer is a very high-dimensionality

search problem, guided by our rapidly increasing knowledge of cancer biology and drugs. The number of potential hypotheses about the causes and associated treatments of the disease is huge, especially when complete genomic profiles and combinatorial therapies are considered—exponentially larger than the number of patients that could represent each possible combination of factors. It is therefore essential to use knowledge to focus the search on promising opportunities, learn as much as possible from every patient, and generalize that knowledge to benefit future patients.

Cancer Commons is an ideal framework for creating human-machine knowledge systems to efficiently conduct this search for better treatments. Think of the networked MDMs as a web-based blackboard and the apps as knowledge sources. The MDMs organize the world's knowledge about cancer subtypes, pathways, tests, drugs, and trials, and that knowledge is made available in a semantic web-compatible format. The human-machine knowledge system that is Cancer Commons generates and tests hypotheses, creates plans, and learns. This shared human-machine approach is essential because although the complexity of the search demands computational management, it is ultimately human medical specialists who interact with patients and who are responsible for carrying out treatment plans. It is imperative that the physicians and patients understand the general rationale behind all decisions and inferences.

The overarching challenges for the AI community are to use the Cancer Commons framework to (1) organize the world's collective knowledge of cancer and use it to predict treatment responses for a given patient; (2) use these predictions to develop treatment plans that provide high expected value for each patient, while simultaneously maximizing the potential for new learning; and (3) learn as much as possible from each patient's molecular and clinical data to improve future predictions while advancing the community's knowledge of cancer. The remainder of this section focuses on AI opportunities and challenges in these three areas.

Knowledge Challenge

Organize the world's collective knowledge of cancer and transform it into personalized, actionable information that can be used to guide treatment decisions.

The molecular disease model approach taken by Cancer Commons posits that a great deal is actually already known about how to treat many genomic subtypes of cancer, but this knowledge is scattered across hundreds of thousands of publications, clinical and biological datasets, and knowledge bases. Medicine has traditionally relied upon the publication of review papers or clinical

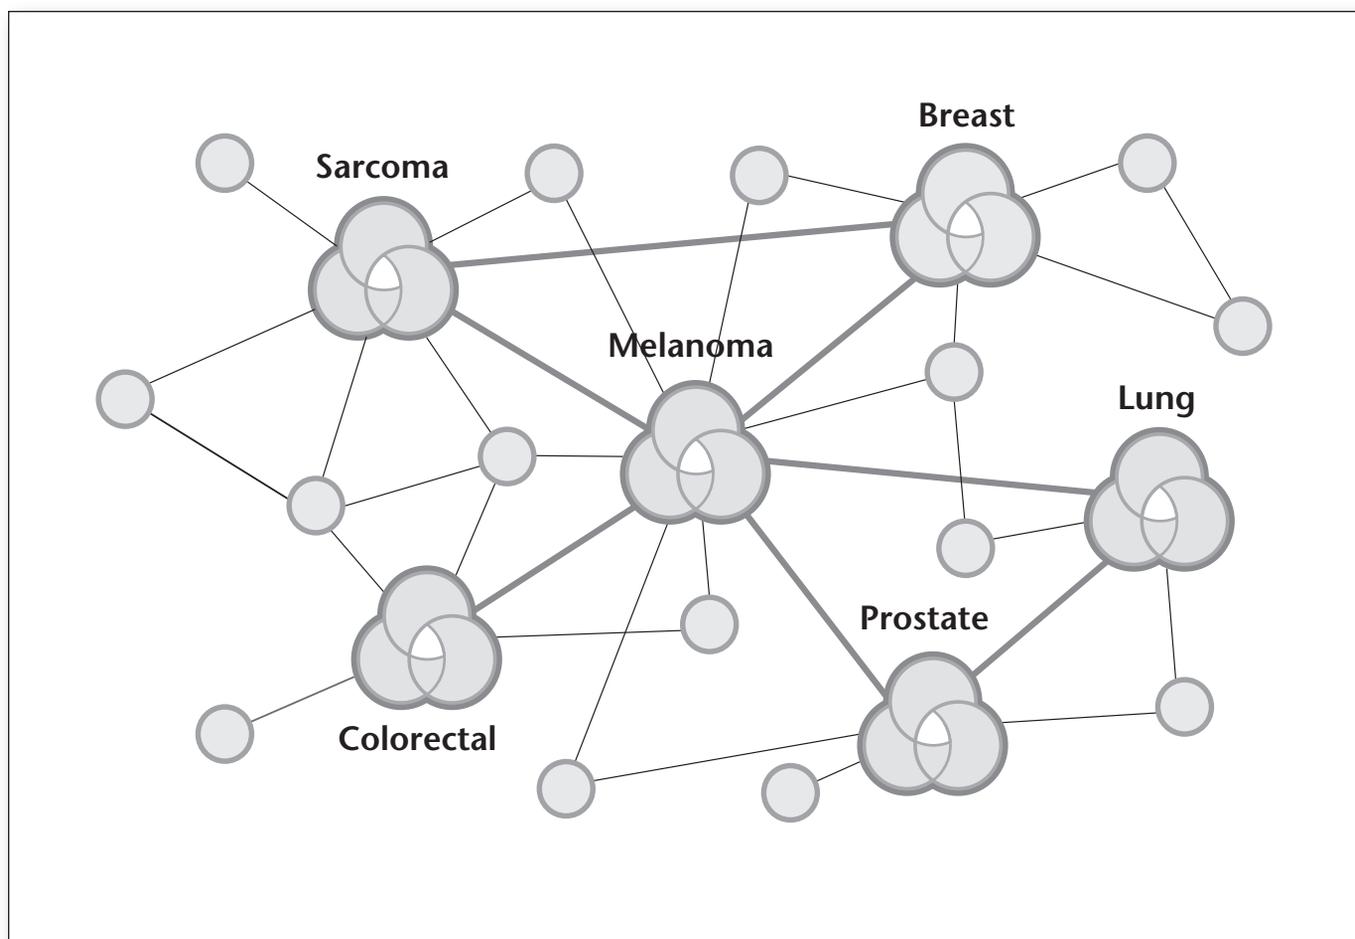


Figure 5. Figurative Depiction of the Cancer Commons Network.

Tripartite icons represent knowledge hubs (for example, MDMs; see figure 3). Circles represent applications and other services that both draw from and feed into the hubs.

guidelines for dissemination of important results, but these compilations have two important limitations: they quickly become dated, especially regarding the latest undocumented results that may only be known to a few experts; and their results are not computationally accessible and so cannot be efficiently utilized.

MDMs are a step forward. As previously noted, they are living review articles, continuously updated by top cancer experts, and their results are available to computers as Resource Description Framework (RDF) triples.¹¹ However it's a safe bet that much of what is currently known about cancer as codified in MDMs is incomplete, inconsistent, uncertain, or downright wrong. Cancer Commons participants can post case reports, laboratory findings, news articles, hypotheses, and the like that support or refute specific MDM hypotheses and recommendations. MDMs can serve as a common point of reference at scientific meetings and in publications to facilitate collaborative research.

Still, organizing the world's knowledge of cancer is easier said than done. The genomics revolution and associated technologies are driving exponential increases in all medical knowledge and data, far outpacing the ability of physicians and patients to keep up. A few examples will illuminate the scale of the problem. PubMed, the definitive medical information retrieval service run by the National Library of Medicine indexes about 20 million medical abstracts. Approximately 750,000 new articles are published annually in medical journals. ClinicalTrials.gov lists about 100,000 active clinical trials. From 80 to 100 trial results are reported daily. The Gene Expression Omnibus (GEO) database includes some 450,000 data sets from gene array experiments. Approximately 25 percent of all these figures involve cancer.

The American Society for Clinical Oncology (ASCO) holds its annual meeting each June in Chicago's McCormick Hall. It is the only venue large enough to accommodate the 35,000 oncolo-

gists and cancer researchers who attend. Over the course of 4 days, some 4000 abstracts are presented, summarizing the results of cancer clinical trials concluded in the past year. About a dozen of those abstracts are deemed important enough to be practice changing and are presented in plenary sessions. One can often read about these in the *New York Times*. The rest effectively never see the light of day. It is likely, however, that many other abstracts, while not practice changing, might be life saving for the right patients; even failed trials often have some responders, just not enough for statistical significance. If only we could figure out what clinical or molecular features characterized these subpopulations and get the lifesaving abstracts in front of those patients and their physicians who might benefit. These are central objectives of Cancer Commons, and AI is essential for achieving them!

Medical advances are also stymied by the serendipity of scientific communications. It can take decades to connect the dots among what is already known to infer a new treatment. The story leading up to the successful testing of Gleevec on c-KIT melanomas actually began in 1991 when Ruth Halaban, a melanoma researcher at Yale, discovered that c-KIT mutations were responsible for the uncontrolled growth of a few percent of melanomas. At the time, there was nothing anyone could do about it. She published a paper in *Cancer Metastasis Review* and moved on (Halaban 1991). In 1998, a team led by Brian Druker at the Oregon Health Sciences University demonstrated that an experimental signal transduction inhibitor known as STI571 was effective in treating CML, a blood cancer caused by a single oncogene (BCR-ABL) (Mauro and Druker 2001). In 2001 Druker and others published an abstract at the ASCO Annual Meeting demonstrating that STI571 was also effective in controlling gastrointestinal stromal tumors expressing mutated c-KIT (Blanke et al. 2001), a protein structurally similar to that produced by BCR-ABL. There was now a promising investigational drug for c-KIT tumors. Three years later, an Italian researcher named Giammaria Fiorentini heard about using STI571 to treat GISTs at an

Italian oncology meeting. He recalled Halaban's 1991 paper associating c-KIT with some melanomas and published a letter to the editor describing a few isolated cases of c-KIT mutated melanomas that they had successfully treated using Gleevec (Fiorentini et al. 2003). Five years later, Stephen Hodi and colleagues at the Dana Farber published a definitive paper in the *Journal of Clinical Oncology* titled "Major Response to Imatinib Mesylate (Gleevec) in KIT-Mutated Melanoma" (Hodi et al. 2008). It took 17 years to connect three dots from Halaban to Hodi, and it will likely take 17 more before this result is widely disseminated to community oncologists.

A human expert like Fiorentini can connect a few dots and experience a "Eureka!" moment. But this only invites the question of how many other effective treatments might be found if one were to systematically (that is, computationally) follow similar connections through the millions of PubMed abstracts and other online knowledge resources. How many subtypes of cancer might already be cured, or at least manageable, if we only knew what we know? AI can help! Table 1 outlines some important AI opportunities in data and knowledge management for Cancer Commons.

Planning Challenge

Adaptively plan individual treatment protocols to achieve optimal outcomes while maximizing the learnings for other patients and cancer research.

From a planning perspective, the fundamental question is: How do you best treat the patient in front of you based on facts about that patient (such as medical history, genomic information), the state of knowledge in the field (results of clinical trials, published facts, and so on), domain expertise (experience of the physicians), and prospective goals (planning of patient's treatment and optimizing learning across the field)? This problem has both retrospective and prospective challenges, the former being classifying the patient and choosing the appropriate therapy, and the latter being the problem of determining what knowledge and data are missing that would help optimize the treatment plan for

the next patient. This analysis must be distributed across all patients. A key challenge is balancing these planning objectives given the inherent tension between treating patients (do everything you can to make this patient better) and populations (learn as much as you can now to help better treat future patients). The planning challenge subdivides into separate challenges in the areas of clinical treatment and research.

Treatment Planning. Treatment planning for an individual is a classic problem of planning under uncertainty and constraints. The basic treatment options could come from standard treatment guidelines, an MDM, or other predictive models. Ideally, each option would be accompanied by a priori likelihoods of success that depend on the molecular subtype of the patient's tumor. Usually, there will also be constraints, including cost, availability (for example, can the patient get into a trial), and general health considerations that might contraindicate a treatment. Sometimes the planner can refine the odds by ordering additional tests, but these carry their own costs and constraints—for example, can the patient afford to defer treatment, are there suitable tumor specimens available for testing.

The planning process must include not just the science, as reflected in today's MDMs, but also qualitative clinical judgments and patient preferences. Even among drugs that "work," there are still judgments that go into decision making that aren't universal. For example, many of today's single-agent targeted therapies have a high likelihood of response in the short term, but the patients are prone to relapse. Immune therapies (for example, Interferon or a cancer vaccine), on the other hand, typically have less than a 20 percent response rate, but those responses tend to be durable. Applying what is known scientifically to individual patients is challenging.

The planning process becomes even more interesting, and (a lot!) more complicated, when it includes combination and sequential therapies. Where more than one drug is involved questions arise concerning dosing and sequencing, for which answers might not be

Modeling
Build or learn models of cancer biology and care that are capable of representing both causal and probabilistic knowledge about diseases, pathways, targets, and drugs.
Use these models to organize the copious information about cancer now scattered across the web for efficient use by people and computers (including data, scientific literature, social media, case reports, electronic medical records, and even speculative hypotheses). Given the state of medical knowledge, the models must be capable of encompassing inconsistent and contradictory information.
Knowledge Acquisition
Populate the models by mining major cancer information sources (for example, journals, websites, blogs, and databases). These miners will identify relevant content and tag it with metadata, thereby linking the content to elements of the disease models (for example, subtypes, pathways, targets, drugs, and trials).
Use domain-specific natural language processing (NLP) to extract knowledge from published content and informal communications (for example, transform a paper on pharmacogenomics into a set of probabilistic assertions about how various mutated genes alter the chemosensitivity of drugs; extract inclusion and exclusion criteria from clinical trial announcements; mine case reports and associated discussions for evidence to support alternative treatment hypotheses).
Search and Question Answering
Use the models to guide deep, vertical search for cancer-related information (for example, find and organize all the information a patient might need in order to evaluate the safety and efficacy of a drug or trial).
Provide direct answers to frequently asked questions (for example, “find the best trials for my cancer subtype” should return a table of ranked candidates).
Inference
Develop domain-specific inference engines that “connect the dots” across knowledge sources to discover new knowledge (for example, discover new uses for existing targeted therapies by searching for subtypes of other cancers that are structurally similar targets; discover new cancer mechanisms by linking two previously disjoint cancer signaling pathways that share a common protein).
Use probabilistic inference to rank targets, drugs, and trials by integrating evidence pro and con from experimental data, the literature, and the Cancer Commons community.
Collective Intelligence
Develop crowd-sourcing approaches that augment automated NLP (for example, community tools for submitting, tagging, ranking, reviewing, commenting on, and recommending content).
Develop tools that make it easy for experts to edit the models and to add or annotate content (for example, suggest new hypotheses and link them to supporting evidence in the models).
Build interactive web forms that enable authors to codify the key concepts in their papers and postings and connect them to concepts in the models. Also, enable them to correct errors introduced by others (for example, NLP and crowd sourcing).
Build tools for collective decision making (for example, enabling physicians to collaborate on a treatment plan that synthesizes their individual experiences and knowledge from the model).

Table 1. Data and Knowledge Management Opportunities.

readily available because the drugs were never tested together in a trial. Indeed, if either of the therapies is still experimental, it may not be available for use in a cocktail. This situation arises frequently with targeted therapies for melanoma and is highly frustrating in the face of compelling evidence that multiple mutated pathways, for which experimental drugs exist, need to be blocked.

Most cancer patients require a sequence of treatments and multiple treatment modalities (such as surgery, radiation, chemotherapy). A treatment may work for a while, but ultimately will have to be augmented or replaced by a second- or third-line therapy when the cancer develops resistance. Good treatment plans maximize “shots on goal.” Picking one treatment can preclude future options (for example, taking drug *A* can exclude a patient from a trial for drug *B*, so to preserve options, drug *B* should be tried first). There is thus a need to consider potential interactions and plan ahead. One should also take into account responses to previous treatments as they can be important clinical predictors of how a patient will respond to a future therapy. Doctors find this kind of optimization difficult, but it is a simple matter for a constraint-

based planner, given the right constraints or the ability to learn them.

Research Planning. Treatment planning for individuals takes place within a larger scientific context. Here the goal is to coordinate planning across all patients to efficiently explore treatment alternatives and fill in knowledge and data gaps. Planning at this level is complicated by the ambiguity and noisy data inherent in medical experiments. It is also strongly constrained by ethics and the interests and incentives of individual physicians.

Balancing treatment and research planning involves the classic exploration-exploitation dilemma from the reinforcement-learning community (Sutton and Barto 1998) but with some important new (ethical) constraints: one would like to do as much exploration as possible while maximizing the outcome for the current patient and respecting that patient’s treatment preferences. It is an important open question as to whether the degree of exploration possible under these constraints will be adequate.

In principle, one could coordinate treatment planning over a population to optimize informa-

<p>Treatment Planning—Develop an incremental planning method that optimizes treatment plans (that is, a sequence of tests, therapies, and trials) for individual patients, based on evolving information and constraints, including:</p> <p>The patient’s clinical history, including diagnoses, treatments and responses.</p> <p>Evolving clinical guidelines for standards of care, and MMDM recommendations (based on the tumor’s subtype) for patients beyond the standards of care.</p> <p>Patient preferences (for example, chemotherapy versus immunotherapy, geographic location, risk tolerance).</p> <p>Availability of drugs (for example, through a trial, or through compassionate use); trial exclusion criteria may dictate the order in which treatments are tried.</p> <p>Availability of tumor specimens for molecular testing (to predict therapy response).</p> <p>Financial considerations, including insurance coverage for tests and drugs.</p> <p>Complicating factors include, for example, dosing and sequencing of combination therapies, prioritizing trials to maximize shots on goal.</p> <p>The planner should support shared decision making by the patient and physician.</p>
<p>Research Planning—Coordinate individual treatment plans to efficiently find the best therapy for each cancer subtype. Overcome ethical conflicts that may arise in treating individual patients by:</p> <p>Randomizing patients to trials or drugs that are deemed equivalent, or whose efficacy is unknown; as evidence mounts, adaptively direct patients toward the more effective therapies.</p> <p>Trying alternative therapies when the best choice is excluded by externalities such as the patient’s state of health, personal preferences, prior treatments, or drug availability (for example, the patient can’t get into a trial because of brain metastases).</p> <p>Utilizing historical controls and observational data to assess trials and drugs that are under-represented in existing Cancer Commons “N of 1” experiments.</p> <p>Globally coordinating the search to ensure adequate coverage of all major drug classes and targets, while minimizing redundant testing of the most popular ones.</p>

Table 2. Planning Opportunities.

tion gain. Imagine, for example, running a definitive set of treatment experiments to systematically rule out an entire family of drugs or drug targets as viable therapeutic candidates, rather than forever revisiting the same dead ends (as is all too often the case today). In practice, it could be hard to convince a group of physicians, with individual interests and incentives, to go along with such a plan. Table 2, summarizing the opportunities for automated planning in Cancer Commons, suggests some ethical strategies for exploiting the daylight between individual benefit and group knowledge acquisition.

Learning Challenge

Use data from thousands of individual treatment experiments to infer the causal mechanisms of tumors and drugs, and develop predictive models for individualizing therapy. Generalize the results so they can be applied to new cases.

Cancer Commons is quintessentially about learning. So far, the focus has been on human learning—tapping the collective intelligence of physicians and patients. However, machine learning can complement human learning by seeking higher-level patterns in the clinical and genomic data that Cancer Commons will collect. Additional insights can be gleaned by mining the massive amounts of drug response data that are becoming available from laboratory experiments.¹²

Knowledge about cancer exists at many other levels of abstraction, such as causal models of tumorigenesis and pathway signaling. Connecting

the dots between these causal models and empirical models of drug response—and filling in the missing knowledge gaps—is hard to do directly. It helps to decompose the overall search into manageable subproblems using intermediate-level models. The molecular disease models at the heart of Cancer Commons are an interesting case in point: MDMs link molecular subtypes of a cancer, which are rooted in causal pathway models, to promising drugs and clinical trials, often selected on the basis of empirical data. MDMs can be used to seed ML algorithms that classify tumors into molecular subtypes (for example, through clustering), as well as algorithms that predict therapy response based on clinical and laboratory data.

There are many other models that are commonly used in cancer research, ranging from the causal models of cancer pathways used in systems biology¹³ to clinical guidelines. The challenge is first to use machine learning to improve these models independently, and then to chain them together to obtain a consistent understanding of cancer at multiple levels. If these challenges can be met, it may one day be possible to use a tumor’s biology to predict drug response in individual patients.

Table 3 proposes some relevant machine-learning projects for Cancer Commons. In addition, there is a rich literature of AI and machine learning applications in cancer biology and medicine that would be interesting to integrate into Cancer Commons. We invite potential collaborators to contact us regarding these and other opportunities.

Cancer Commons provides a rich context for

"N of 1" Experiments—Given clinical and genomic data from a series of patients with a given type of cancer.
Apply hierarchical clustering techniques to determine subtypes and validate them against those in the expert-curated MDM.
Determine the optimal therapy for each individual, based on previous clinical outcomes and therapy responses for patients with the same subtype.
Replace clustering with more sophisticated approaches to tumor subtyping, such as using genomewide expression and sequence data to infer dysfunctional upstream pathways and then clustering patients with the same aberrant pathways.
Replace subtyping with a similarity metric that matches patients directly, based on their molecular tumor profiles and responses to previous therapies; use this metric to recommend future therapies, based on what worked in similar patients.
Clinical Trials—Reanalyze clinical and molecular data on individual patients in population-based trials.
Analyze the biological differences between responders and nonresponders to discover predictive biomarkers. Use them to salvage drugs that were effective in subpopulations. Meta-analyze data from multiple trials to enhance statistical power.
Design "intelligent" adaptive Bayesian trials ¹⁴ that use AI techniques such as reinforcement learning (RL) to efficiently test molecular diagnostic and targeted therapy candidates. Unlike traditional trials, which test one or two drugs on a population, the goal here is to predict responses and maximize outcomes for individuals for a continuously evolving set of drug and biomarker candidates.
Laboratory Studies—Given results of high throughput drug screens on cancer cell lines.
Build a predictive model of drug response based on the molecular profiles of the cells (for example, from microarray data). ¹⁵
Use the model to recommend therapies for patients, based on the drugs that worked best for genomically similar tumor cells in the laboratory. If resources permit, validate the prediction in vitro or in vivo as appropriate.
Use laboratory and clinical results to continually refine the predictive model.
Literature Studies—Given a corpus of PubMed articles reporting on various interactions among genes, proteins, and drugs.
Infer network (pathway) models of cancer biology by integrating fragments of evidence on gene-protein and protein-protein interactions from the literature or from actual experiments (for example, Bonetta [2010]).
Meta-analyze the extensive literature on pharmacogenomics (the effect of genes and gene expression on drug response and drug-drug interactions) to create predictive models for drugs (that is, what are the optimal drugs for a patient whose tumor cells express particular genes and proteins) (for example, Mocellin et al. [2010] and Theobald, Shah, and Shrager [2009]).
Test predictions in the laboratory and ultimately in patients.

Table 3. Potential Machine-Learning Projects for Cancer Commons.

ML projects like those proposed in table 3. There will be opportunities to cross-validate results from human and machine learning, as well as from causal and empirical models. There will be "transfer learning" opportunities across different organ-based cancers through the Cancer Commons network—using an approved cancer drug off label to treat a different cancer with a similar molecular tumor profile. The goal is to generalize the findings from a few patients to a much larger group of patients like them. Opportunities like these will drive research in both ML and cancer and, we hope, save lives.

Conclusion

Cancer Commons is the harbinger of a new patient-centered paradigm for cancer research in which every patient receives individualized therapy based upon the best available knowledge and where the resulting data contributes as efficiently as possible to improved treatment for each subsequent patient. AI and machine learning will be essential to realizing this vision.

Cancer, as a domain, presents the AI and ML communities with extraordinary challenges and, hence, opportunities. These include (1) managing the sheer quantity and heterogeneity of the data and knowledge involved—encompassing millions of medical records, genomewide datasets, and documents; (2) planning thousands of complex, multistep treatment strategies that ethically balance

the needs of the individual with those of science; (3) capturing and analyzing the results of these treatment experiments in diverse causal and empirical models of cancer biology and drug response; (4) continuously testing and refining these models to account for new clinical and laboratory findings; (5) generalizing the models across patients and cancers and integrating them to improve decision making; and (6) integrating human and machine planning, learning, and decision making to exploit their respective strengths.

Taken together, these challenges add up to the grand challenge of curing cancer. This challenge is a worthy heir to the succession of grand challenges—from chess to unmanned vehicles—that have driven so much progress in AI over the years. While AI may not be able to cure cancer any time soon, it may soon make life or death differences in outcomes for individual patients. Please join us and the Cancer Commons community in rising to this challenge.

Grand Challenge

Develop AI applications that tap the world's collective knowledge and data on cancer to improve patient outcomes and save lives.

To learn more about this grand challenge, please visit www.cancercommons.org.

Acknowledgements

The authors are indebted to our colleagues at Col-labRx and on the Cancer Commons melanoma

editorial board, especially Smruti Vidwans, Allan Schiffman, David Fisher, MD, and Keith Flaherty, MD, for their contributions to the melanoma molecular disease model and therapy finder application. We also acknowledge the thoughtful input we received from colleagues in machine learning and computational biology, particularly Atul Butte, David Haussler, and Daphne Kohler who compose the Cancer Commons computational advisory board, and Michael Littman and Sriram Natarajan who helped organize a workshop and team proposal on AI and cancer.

Notes

1. See www.cancer.gov/cancertopics/factsheet/Therapy/targeted.
2. See www.alternet.org/health/140234.
3. See www.CancerCommons.org.
4. "Translational research" is defined as the translation of scientific discoveries into practical applications (commonfund.nih.gov/clinicalresearch/overview-translational.aspx).
5. See www.economist.com/node/14299624.
6. The MMDM was curated by a panel of 11 melanoma experts led by cochief editors David Fisher, MD, head of dermatology, and Keith Flaherty, MD, head of experimental therapeutics, at MGH.
7. See mmdm.cancercommons.org/smw/index.php/A_Melanoma_Molecular_Disease_Model.
8. See semantic-mediawiki.org.
9. See collabrx.com.
10. See therapy.collabrx.com.
11. Resource Description Framework (RDF) triples are used in the semantic web to encode knowledge in the form of subject-predicate-object expressions.
12. See, for example, sagebase.org.
13. See, for example, Ingenuity.com.
14. See, for example, www.sciencenews.org/view/generic/id/58401/title/BATTLE_trial_personalizes_lung_cancer_treatment, <http://www.ispy2.org>.
15. See www.broadinstitute.org/cmap.

References

- Blanke, C. D.; von Mehren, M.; Joensuu, H.; Roberts, P. J.; Eisenberg, B.; Heinrich, M.; Druker, B.; Tuveson, D.; Dimitrijevic, S.; Silberman, S. L.; and Demetri, G. D. 2001. Evaluation of the Safety and Efficacy of an Oral Molecularly-Targeted Therapy, STI571, in Patients (Pts) with Unresectable or Metastatic Gastrointestinal Stromal Tumors (GISTs) Expressing c-KIT (CD117). *Proceedings of the American Society of Clinical Oncology* 20: (abstr 1). Alexandria, VA: American Society of Clinical Oncology.
- Bonetta, L. 2010. Protein-Protein Interactions: Interactome Under Construction. *Nature* 468: 851–854.
- Demetri, G. D. 2002. Targeting the Molecular Pathophysiology of Gastrointestinal Stromal Tumors with Imatinib. Mechanisms, Successes, and Challenges to Rational Drug Development. *Hematology Oncology Clinics of North America* 16(5): 1115–24.
- Fiorentini, G.; Rossi S.; Lanzanova G.; Bernardeschi, P.; Dentico, P.; and De Giorgi, U. 2003. Potential Use of Imatinib Mesylate in Ocular Melanoma and Liposarcoma Expressing Immunohistochemical c-KIT (CD117). *Annals of Oncology* 14(5): 805.
- Grove, A. 2005. Efficiency in the Health Care Industries: A View from the Outside. *Journal of the American Medical Association (JAMA)* 294(4): 490–492.
- Haliban, R. 1991. Growth Factors and Tyrosine Protein Kinases in Normal and Malignant Melanocytes. *Cancer Metastasis Review* 10(2): 129–40.
- Hodi, F. S.; Friedlander P.; Corless, C. L.; Heinrich, M. C.; Mac Rae, S.; Kruse, A.; Jagannathan, J.; Van den Abbeele, A. D.; Velazquez, E. F.; Demetri, G. D.; and Fisher, D. E. 2008. Major Response to Imatinib Mesylate in KIT-Mutated Melanoma. *Journal of Clinical Oncology* 26(12): 2046–51.
- Kim, K. B.; Eton, O.; Davis, D. W.; Frazier, M. L.; McConkey, D. J.; Diwan, A. H.; Papadopoulos, N. E.; Bedikian, A. Y.; Camacho, L. H.; Ross, M. I.; Cormier, J. N.; Gershenwald, J. E.; Lee, J. E.; Mansfield, P. F.; Billings, L. A.; Ng, C. S.; Charnsangavej, C.; Bar-Eli, M.; Johnson, M. M.; Murgo, A. J.; and Prieto, V. G. 2008. Phase II Trial of Imatinib Mesylate in Patients with Metastatic Melanoma. *British Journal of Cancer* 99(5): 734–740.
- Mauro, M. J., and Druker, B. J. 2001. STI571: Targeting BCR-ABL as Therapy for CML. *The Oncologist* 6(3): 233–238.
- McDermott, U.; Downing, J. R.; Stratton, M. R. 2011. Genomics and the Continuum of Cancer Care. *New England Journal of Medicine* 364(4): 340–350.
- Mocellin, S.; Shrager, J.; Scolyer, R.; Pasquali, S.; Verdi, D.; Marincola, F. M.; Briarava, M.; Gobbel, R.; Rossi, R.; and Nitti, D. 2010. Targeted Therapy Database (TTD): A Model to Match Patient's Molecular Profile with Current Knowledge on Cancer Biology. *PLoS ONE* 5(8): e11965. doi:10.1371/journal.pone.0011965
- Pleasant, E. D.; Cheatham, R. K.; Stephens, P. J.; McBride, D. J.; Humphray, S. J.; Greenman, C. D.; Varela, I.; Lin, M.-L.; Ordóñez, G. R.; Bignell, G. R.; Ye, K.; Alipaz, J.; Bauer, M. J.; Beare, D.; Butler, A.; Carter, R. J.; Chen, L.; Cox, A. J.; Edkins, S.; Kokko-Gonzales, P. I.; Gormley, N. A.; Grocock, R. J.; Haudenschild, C. D.; Hims, M. M.; James, T.; Jia, M.; Kingsbury, Z.; Leroy, C.; Marshall, J.; Menzies, A.; Mudie, L. J.; Ning, Z.; Royce, T.; Schulz-Trieglaff, O. B.; Spiridou, A.; Stebbings, L. A.; Szajkowski, L.; Teague, J.; Williamson, D.; Chin, L.; Ross, M. T.; Campbell, P. J.; Bentley, D. R.; Futreal, P. A.; and Stratton, M. R. 2010. A Comprehensive Catalogue of Somatic Mutations from a Human Cancer Genome. *Nature* 463 (14 January): 191–196.
- Shrager, J.; Tenebaum, J. M.; and Travers, M. 2011. Cancer Commons: Biomedicine in the Internet Age. In *Collaborative Computational Technologies for Biomedical Research*, ed. by S. Elkins, M. Hupcey, and A. Williams. New York: Wiley.
- Sutton, R., and Barto, A. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: The MIT Press.
- Theobald, M.; Shah, N. H.; and Shrager J. 2009. Extraction of Conditional Probabilities of the Relationships between Drugs, Diseases, and Genes from PubMed Guided by Relationships in PharmGKB. Paper presented at the AMIA Summit on Translational Bioinformatics, San Francisco, March 15–17.
- Vidwans, S.; Flaherty, K. T.; Fisher, D. E.; Tenebaum, J. M.; Travers, M. D., and Shrager, J. 2011. A Melanoma Molecular Disease Model. *PLoS ONE* 30 6(3): e18257. doi:10.1371/journal.pone.0018257.
- Jay M. "Marty" Tenenbaum**, Ph.D., was educated at the Massachusetts Institute of Technology and Stanford University in the 1960s. He spent the 1970s doing artificial intelligence research at SRI, the 1980s managing computer science research for Schlumberger, and the 1990s pioneering Internet commerce. He's currently focused on using AI and the web to transform medicine.
- Jeff Shrager**, Ph.D., is the CTO of CollabRx and a consulting associate professor in the Symbolic Systems program at Stanford University. His research focuses on understanding how science works and how computers can facilitate scientific discovery. Shrager has cofounded several scientific-computing startups and is the inventor of BioBike, a web-based biological knowledge platform enabling biologists to develop, run, and share symbolic analyses of genomic information.