ELSEVIER

feature

# To measure is to know: an approach to CADD performance metrics

**Deborah Loughney, Brian L. Claus and Stephen R. Johnson,** stephen.johnson@bms.com

**With the financial and productivity challenges currently facing the pharmaceutical industry, there is constant pressure to justify resources and improve efficiency. With process-driven activities, understanding the contribution of these resources is reasonably straightforward. By contrast, measuring the contribution of knowledge workers is less obvious. Here, we present an impact-oriented approach to assessing the performance of an industrial computer-assisted drug design group. We discuss how these metrics are used to understand and optimize resource allocation in support of drug discovery programs.**



## Introduction

The issues affecting the pharmaceutical industry are well documented, from the sharp loss of product exclusivity to the decline of productivity in bringing new drugs to market [1]. As a direct result of these issues, research groups are under pressure industry-wide to do more with fewer or equivalent resources. This, in turn, has driven the need to better optimize how staff work within pharmaceutical R&D and to better deploy the resources currently available. It is difficult, if not impossible, to optimize reliably the performance of a group without a good understanding of its current performance.

Metrics themselves are nothing new to the business community, having found widespread application in nearly every industry conceivable, including pharmaceuticals [2,3]. Still, the larger question of how (or even if) to apply them to creative scientific work lingers when discussing the topic with peers. There is a perception among scientists in general that their work is 'beyond metrics' or that metrics serve only as a barrier to creative work and scientific progress. To be sure, incorrect metrics can do great harm. In the landmark paper 'On the folly of rewarding for A, while hoping for B' [4], Kerr discusses several case studies of harmful outcomes of misguided metrics. Still, a central tenet of that article was that correct metrics would indeed yield positive outcomes. Stated more simply, you get what you measure.

The question of metrics in science is one of interest in the literature, with a series of articles appearing recently regarding the potential assessment of academic research by granting organizations [5,6]. Most of these approaches use citation data from the scientific literature as a surrogate for research quality and importance [7]. Arguably, although this approach might be reasonable for assessing academic research effort, it is not a tractable approach for evaluating the utility of commercial scientific research. This is largely because most research in industrial laboratories is not published in the primary literature for intellectual property reasons. Even excluding this reason, a particular contribution could be important for a relatively small audience, but would be given low scores using citation-based assessments.

As an alternative, we describe here a series of outcome-oriented metrics to track and assess the performance of the computer-assisted drug design (CADD) group at Bristol-Myers Squibb (BMS). Generally, the primary objective of a CADD scientist is to provide hypotheses founded on sound scientific principles that lead to molecules and proteins with improved pharmaceutical properties and activities. The primary CADD deliverable is an idea. Rather than counting the number of ideas generated, we focus on the overall impact that these ideas have on the downstream research effort. One important use of metrics, if not the most important one, is to provide indicators as to the strengths and weaknesses of the group to allow for further improvement. We discuss several such efforts that were shaped by these metrics.

☆ The title is a play on a quote by Lord Kelvin: 'To measure is to know'.



1359-6446/06/$ - see front matter  doi:10.1016/j.drudis.2011.05.003

**BOX 1**

### Knowledge work

Peter Drucker is usually credited with the first use of the term 'knowledge worker' [13]. A knowledge worker is someone with specific domain knowledge who specializes in separating relevant information from irrelevant information. Drucker has since given six keys to knowledge worker productivity [14]:

(i) Defining the task: in contrast to manual work, the task that a knowledge worker needs to perform to solve a problem is not obvious and must be determined by the worker themselves.
(ii) Knowledge workers must be autonomous: they must be responsible for their own productivity.
(iii) Innovation must be central to the work and be the responsibility of the worker.
(iv) Knowledge workers must be equally continually learning and continually teaching.
(v) Productivity is not about quantity of output. 'Quality is the essence of the output.' [14]
(vi) Knowledge works must be considered an 'asset' and not a 'cost'.

As a result of these points, Drucker stresses that the best approach to knowledge worker productivity is to stress quality over quantity. This, in turn, means that we must define quality. In our approach, we define quality by the impact that the work has on downstream research efforts.

### Evolution of metrics for CADD

At BMS, the CADD group is part of a larger technology organization, which includes functions such as high and medium throughput screening groups, compound management, laboratory automation and structural biology. The grouping of these seemingly diverse functions in one organization creates synergies that ultimately lead to higher impacts on programs owing to the input of cross-functional teams on the discovery process. Many of these teams are process-oriented technology groups and the nature of their work lends itself to metrics that are used effectively to drive performance and efficiency. By contrast, CADD research is primarily 'knowledge work' (Box 1), which is more difficult to quantify and analyze.

The BMS CADD group comprises individuals with diverse backgrounds and skills sets, resulting in creative approaches to problem solving. The notion of measuring any part of the work process was a foreign one and was met with some resistance (Box 2). This negative reaction can be understood in the context of the advice given by Drucker [8], who stated that the best approach to managing knowledge workers is one where management should 'strip away anything that gets in their way'. So where does one start to develop metrics for a discipline that has no apparent process and delivers intangibles? One could, for example, easily envision counting central processing unit (CPU) minutes consumed, but to what end? That begged the question of how other industries where highly creative individuals are employed measure their progress and productivity. Indeed, the need for different metrics for creativity-driven activities was noted previously [9]. It is clear from a survey of the literature that knowledge workers are the focus of much interest in the arena of performance management.

**BOX 2**

### Resistance to metrics: quotes from CADD scientists

When we first attempted to define performance metrics for CADD, there was substantial resistance from members of the group. One of the most valued attributes of the group is the freedom to contribute feedback candidly before decisions are made. Few topics push the boundary of this freedom further than metrics did. This is just a sample of quotes taken from the initial e-mail thread introducing the topic to the research group members:

- 'It's obvious to anyone with half a melon in their head that applying a number to the work we do is complete and utter nonsense'.
- 'When will we be able to stand up and simply say that CADD is exempt from such meaningless metrics?'
- 'Metricizing is a concept most applicable to measuring factory output, sales, and dieting. It is concept foreign to the creative process intrinsic to CADD. The work done within CADD cannot be measured.'
- 'This fixation on metrics is dangerous.'
- "If the group believes strongly that metrics do not apply, then I would think that a Ghandi-esque approach would work quite well, even if we are 'voices crying in the desert.'"
- 'Spending time to create yet more abstract measures of performance in an attempt to industrialize an inherently unindustrializable process does little but distract us from our ultimate goal of discovering drugs.'

### Approach to metrics

One criticism that came up repeatedly in our conversations about metrics in science is that tracking tasks in science is not useful. Clearly, this is especially true in computational modeling, where the number of calculations that can be performed relatively easily is essentially unbounded. This opinion was only further reinforced by the literature on metrics for creative work typically performed by knowledge workers [10,11], which stressed the importance of measuring impact as a measure of quality rather than counting tasks.

Our first step on the road to establishing metrics for CADD was to define what we do and how we do it. The CADD process was defined and reduced to three basic components (Fig. 1): data gathering, data analysis and scientific decision support. Data gathering is self-explanatory and includes chemical, biological, biophysical or any other data that can be brought to bear on a problem. The analysis phase is where computational chemistry tools and techniques are applied to reduce data and information to knowledge by the application of expert opinion, skill sets and experience. Scientific decision support is where the CADD scientist provides guidance to his or her collaborators as to what should be done to test a hypothesis. The deliverable is an idea or advice that has no intrinsic value until it is applied or acted upon in some way by another member of a drug discovery program team, making the final output from this process flow an impact on a program. These three components provided context for the decision to use impacts as our primary metrics tool.

We then sought to define our arena of impact. This began with a goal-oriented mission statement culminating in a relatively simple statement: to guide the identification, selection and optimization of drug candidates and targets

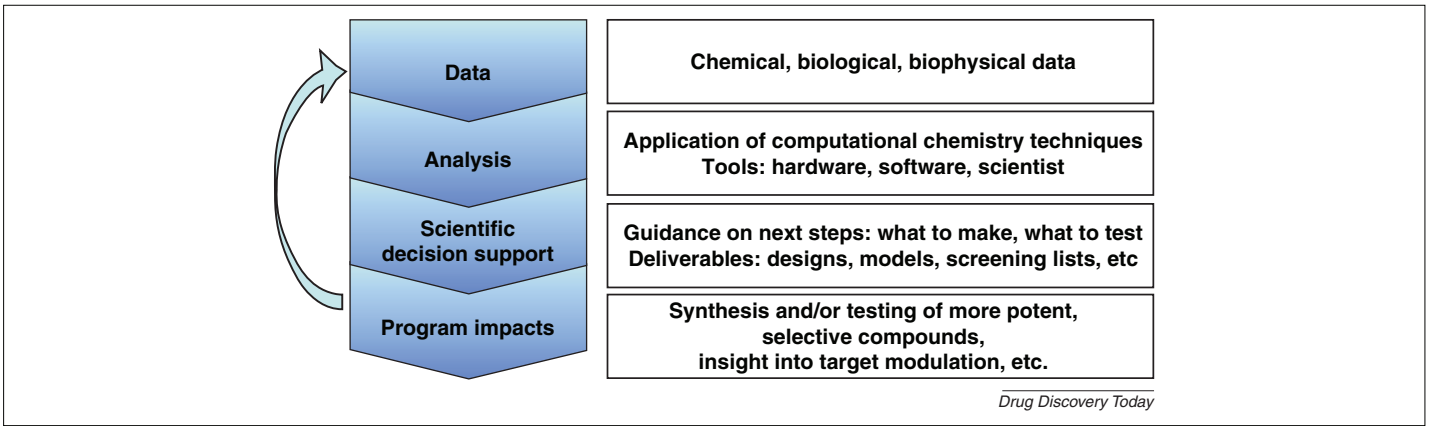


**FIGURE 1**

General process outlined for CADD contributions to projects. The process begins with collecting the relevant available data, analyzing the data using computational tools and then providing guidance to the project teams.



through the application and development of molecular modeling, data analysis and computational chemistry techniques.

This mission statement highlights both what we are going to do and how we are going to do it. The metrics flow directly from the mission statement and are intended to track how well we accomplish that mission. We want to measure our ability to identify, select and optimize compounds and targets; these activities roughly correspond with the discovery component of the Product Development and Commercialization (PD&C) process used at BMS. We also want to measure the impact of our internal method development efforts.

To track impact, the contributions of CADD to projects are categorized as either data provided, significant or enabling. The definitions of these categories are qualitatively similar to those used by Houston *et al.* [12] but have evolved to include impacts on programs by all the functions in our organization. The lowest category, 'data provided', simply reflects that some information or hypothesis was transferred from CADD to another scientist on a project. A contribution is considered 'significant' if the receiving scientist acts on the information or hypothesis and the experimental outcome is consistent with contribution of CADD. Finally, the highest level of impact, 'enabling', is achievable only when a project leader agrees in writing that a contribution substantially altered the direction or timing of the project. This category is difficult to achieve intentionally.

As an example, a CADD scientist constructs a homology model of a protein and uses docking to propose a potential binding mode of a known active ligand. He or she presents this model to the chemists on the project and suggests a compound with a novel core that would maintain similar interactions with the protein. At this point, the binding model and proposed compound are considered to be 'data provided' to the project. A chemist on the project later synthesizes the proposed compound and it is found to have similar or improved activity relative to the known active. The contribution could then be upgraded to significant. After several tens or even hundreds of compounds containing this core are synthesized, the project leader agrees that the contribution enabled the success of the project as the novel chemotype solved a particularly difficult and/or important problem.

## The burden of metrics

As the group discussed the objective of defining metrics for CADD, the only area of initial agreement was the desire to spend as little time collecting metrics as possible. It would, after all, be somewhat foolhardy to spend more time tracking and calculating metrics than performing the actual work being measured. In response to this, we have automated as much of the metrics process as possible. Nonetheless, there is always some amount of data entry that must be manually maintained.

Central to the metrics calculations are two databases, both of which pre-dated our metrics effort. The first database is used to track the time spent on projects by way of time sheets filled in monthly by every scientist. This database is mandated by upper management to allow them to track and manage the project portfolio. Given that the database is multipurpose, there is no additional burden from the scientist's perspective to make this information available for the metrics process.

The second database is used to track the impacts on each project. The impact database is maintained on a monthly basis by the group leaders to ensure accurate tracking of the role of each group with respect to the progress of a project. Functionally, every scientist in the CADD group is expected to produce a monthly list of key program contributions. This report generally reads similar to an expanded table of contents from the scientists' laboratory notebook. The group leaders take these reports and input the material into the impact database, categorizing each piece of information into one of the three impact bins.

The information from these databases is extracted programmatically and collated, pivoted and reformatted as necessary to generate the metrics. This information is merged with portfolio information so that the stage of individual projects (exploratory, early no chemistry, early with chemistry, full phase and backup), target class, therapeutic area, and so on, can be identified. In general, the resulting metrics are aggregated by the project rather than by the scientist. Therefore, if two or more scientists contribute to a single project, the time and impacts from all of the scientists are combined.

One important caveat to consider is why a given contribution has no impact on a program. There are several possible causes for this. One is that the CADD support was inadequate or unnecessary. This is obviously the type of result that we would like to avoid. Another reason might be the temporal nature of waiting for an outcome. As we are measuring impact, most CADD analyses begin as a 'data provided' outcome, only to be upgraded later if they are acted upon. Therefore, it is possible that many of the

non-impacted projects simply had not matured or were terminated before anyone acting on the CADD results.

### Five key questions

Our intention was that metrics should flow naturally from the mission statement and be outcome rather than input oriented. This focus on outcomes, rather than tasks, is an important component of the approach to metrics taken here. In contrast to manufacturing, where a physical entity is created or assembled, the direct result of CADD modeling is typically an idea, hypothesis, or proposed experiment. Each of these has value, of course, but assessing the value of a particular idea, hypothesis, or experiment is nontrivial. Perhaps the easiest way to assess that value is by judging its impact *post facto*. Effectively, we are using the follow-up on a CADD proposal as a surrogate for the quality of the ideas generated and for the influence that CADD has within the project teams.

In that vein, our metrics were aimed at answering five key questions regarding impact: (i) what is our capacity to support discovery projects? (ii) What percentage of supported projects do we impact? (iii) Where on the discovery timeline is CADD most impactful? (iv) How impactful is innovation to our work? And (v) how does our resource allocation effect our impact?

#### *What is the capacity of CADD to support projects?*

At BMS, a CADD scientist is typically assigned to a discovery project to support its molecular modeling needs. We are constantly being asked to staff additional projects or initiatives. To address the question of the overall capacity of the CADD group, we compared, using data aggregated over the trailing 3 years, the median full-time equivalent (FTE) time spent on projects with at least a significant or enabling outcome, to the median FTE time spent on projects for which the highest impact level was merely 'data provided' or 'no contribution'. Fig. 2 shows these data for different stages of the discovery timeline.

Projects that are staffed more heavily tend toward more impactful contributions by CADD. For example, a median of approximately 0.4 FTEs are invested in full-phase primary projects with an impactful outcome, compared with only 0.12 FTE for projects with no impact. Based on this impact data, a typical CADD scientist should be able to support two to three projects before being spread too thin to make a meaningful contribution.

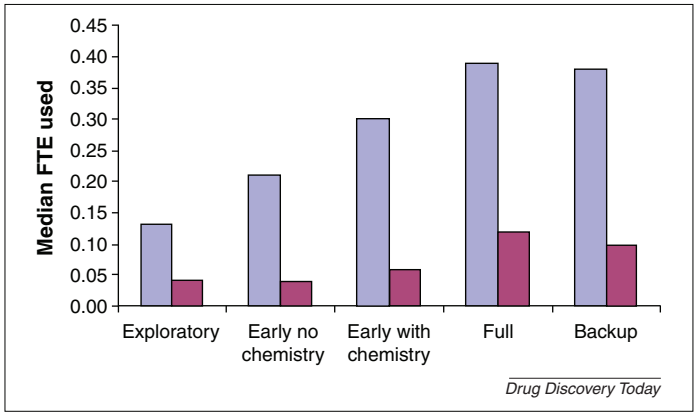


**FIGURE 2**

Median FTE spent on projects with significant or enabling impacts (blue bars) compared with projects with no impact or data provided (maroon bars). Values shown are results aggregated from 2007 to 2009.

This information led to a more definitive description of the capacity of the group than was previously possible. When asked about the CADD resources required to support a new initiative, for example, traditional surrogates, such as the ratio of CADD scientists to chemists, were not always appropriate. We believe that basing our capacity estimates on impact data provides a more realistic projection of the bearable load that our group can productively sustain and validates the use of projects per person as the appropriate capacity metric. It also provides a quantity that we can track over time to assess the effect of technology improvements and innovation on our ability to support projects.

The second outcome of this metric was the development of a more definitive approach to effective program support. There is a clear trend in our data that shows that lightly resourced projects are only marginally impacted by CADD. Therefore, we now actively avoid these scenarios, choosing instead to focus larger fractions of our time on fewer projects. There are exceptions to this general guideline, wherein low time commitment tasks have large potential for impact. The metrics do not preclude us from working on these but are rather a tool to guide the prioritization of activities. Metrics are a useful tool for these conversations but cannot be used in a vacuum to make blanket decisions on resourcing.

#### *What percentage of projects do we impact?*

One goal of our organization is to make as broad an impact on the project portfolio as possible. Our metric to address this question is simply the percentage of projects with a maximum impact of 'no contribution', 'data provided', 'significant' or (enabling). These data are shown in Fig. 3. No contribution assignments occur when a project has time reported for a CADD scientist but no obvious contribution recorded for the project. Gains in significant and enabling impacts in recent years have come from several factors, including better alignment with project focus, the elimination of 'drive-by' CADD work, better reporting and innovation.

#### *Where on the discovery timeline is CADD most impactful?*

CADD provides project support from the very beginning of the target identification process up to, and including, the nomination of a backup compound for preclinical workup leading to First in Human trials. Fig. 4 shows the impact/FTE of CADD scientists by project status along our PD&C process. Over the past 4 years, CADD productivity has been steadily improving in all stages of the discovery timeline, except for support of backup programs, which has been somewhat variable.

Frequently after the nomination of a compound from a discovery project, a decision is made to pursue at least one additional compound to address any perceived concerns regarding the primary compound, and this is considered a separate backup program. Backup programs can generally be divided into two types, a fast follow-on or a novel chemotype program. The novel chemotype backup program has impact effectiveness that is roughly equivalent to early-phase programs, as the goal of the program is essentially the same. By contrast, a fast follow-on project is typically searching for a compound closely related in structure to the primary molecule. Our analysis shows that, despite a roughly equivalent number of data-provided impacts, fewer convert to significant or enabling owing to the nature of the

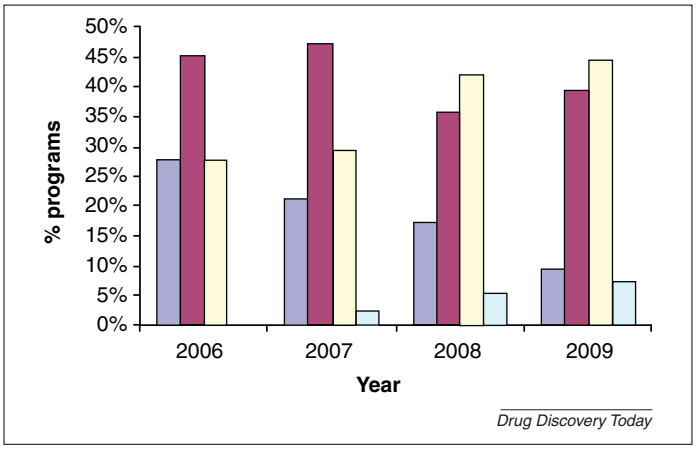


**FIGURE 3**

Percentage of discovery programs by highest impact classification for each year. Key: blue bars, no contribution; maroon bars, data provided; yellow bars, significant impact; aqua bars, enabling.

problems being addressed. Backup projects typically work within a more constrained chemistry on problems that are less amenable to structure-based design or other well-developed modeling methodologies. Partially as a result of this analysis, CADD has adopted a strategic plan to address weaknesses in these types of end-point, including a renewed focus on Absorption, Distribution, Metabolism, Excretion, and Toxicology (ADMET) modeling and multiobjective design tools for designing compounds.

### *How impactful is innovation to our work?*

In any scientific discipline, innovation is the lifeblood of continued relevance. Within the CADD field, innovation largely takes the form of software development for the purposes of novel methodology implementation. Broadly speaking, our methodology development activities arise from two sources. The first arises during the course of project support. As our approach to evaluation is outcome oriented, scientists are free to tackle problems in the manner they consider most appropriate, potentially resulting in the implementation of novel techniques. Second, a more traditional gap analysis focuses internal innovation efforts on key activities that are central to the CADD mission. As part of that analysis, the allocation of resource is weighted against the potential for impact. One constant challenge for internal development is finding the correct balance between rapid prototyping versus a formalized software development lifecycle. By having a distinct innovation metric, there is incentive to evolve the tools to a point where they are accessible to the rest of the group. Guidelines for common implementation frameworks and documentation ensure reusability and accessibility beyond the developer.

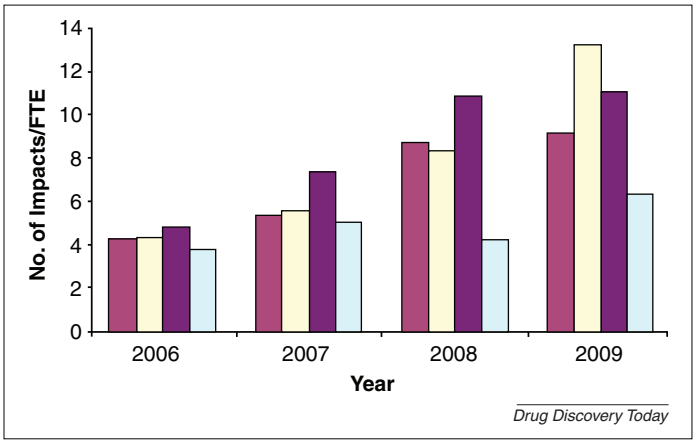


**FIGURE 4**

Impacts/FTE by project status. Impacts/FTE are calculated by determining the number of (Significant + Enabling Impacts)/(Number of hours/2040). Key: maroon bars, Early no chemistry; yellow bars, Early with chemistry; purple bars, Full phase (primary); aqua bars, Full phase (backup).

Fig. 5 shows the number of impacts from projects leveraging internal innovation. Although some of these innovations could be construed as 'workflow' tools, most are not; neither are they ground-breaking Nobel Prize-worthy scientific accomplishments. In general, they are small scientific innovations around tasks such as bioisostere identification, library design, molecular similarity and alignment, hydrophobicity or Structure-Activity Relationship (SAR) analysis. Fig. 5 also shows the number of innovative methods that were used in generating the impacts. Although growing, these numbers reflect our recent focus on small numbers of efforts targeted at improving particular aspects of our scientific capabilities.

### *How does our resource allocation affect our impact?*

The foremost annual expense of a CADD group, excluding salaries and other personnel expenses, is the annual budget dedicated to licensing software. In response to a resource-constrained environment coupled with increasing pressure to improve productivity, we developed and executed a strategy to better leverage our software budget. We used our impact-oriented approach to metrics to track progress of the strategy over time and to guide resource allocation. Fig. 6 shows the relationship between project impact and the distribution of our software budget. In 2004, the group made a conscious decision to do two things. First, we sought to simplify our compute environment by reducing the number of software vendors we engaged to do our work and by standardizing the compute hardware used. This reduced the effort dedicated to software maintenance, such as installations, custom operating system (OS) environment setups, such as graphics card support, license server upkeep, and so on, and allowed us to expand compute capacity in a cost-constrained environment. Second, we adopted a single modeling platform that removed non-value-added work by standardizing common tasks, simplifying communication within CADD and facilitating dissemination of results to our drug discovery colleagues.

In 2006, a focus on internal innovation was added as a third strategic priority to the software portfolio. The selection of vendor A as the standard modeling platform, and the standardization of the compute hardware used,

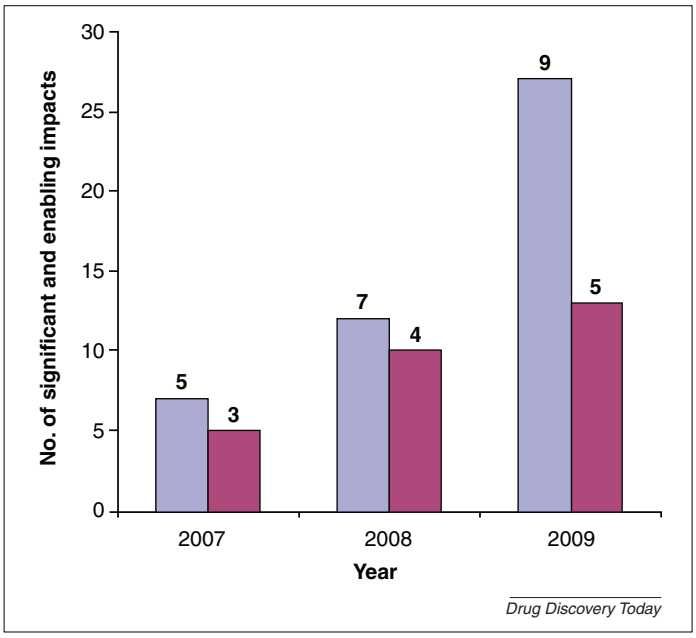


**FIGURE 5**

Impacts derived using methodology implemented in-house. Methods are divided into those conceived internally (blue bars) versus those implemented based on literature reports (maroon bars). The values above each bar represent the number of unique methods leading to the impacts reported.

simplified the compute environment, allowing a rededication of internal effort and funds within the CADD group to focus on developing innovative scientific methods. This trend is reflected in the increasing commitments to vendor B as the underlying platform for unifying internal methodology development efforts. It is clear from Fig. 6 that this combination of a simplified modeling environment with project-focused methods development has resulted in an increase in impacts per FTE while maintaining a flat budget.

### Metrics in performance management

Other than solely being a tool to communicate to upper management how the group is faring, one of the tenets we held was that the metrics should be useful for directing continuous improvement efforts. This outcome-oriented approach to assessing our work is now an active part of the objective setting process for both groups and individuals. Individual scientist objectives state that they should have a minimum number of impacts on any program that they support. The number of impacts varies with the expected difficulty of the program. This allows early intervention as projects are observed to have used substantial FTE time from CADD with little or no resulting impact. By looking across the two or three projects that a CADD scientist is supporting at a given time, we can begin to understand whether low impact is project, scientist, or time related. This analysis can then prompt additional analysis or other intervention as needed. It would be dangerous to use these metrics in isolation. The metrics are used in conjunction with oral and written feedback from the project and CADD leadership to ensure that the scientist is focused on the relevant issues and not just on the simple to-do tasks.

### Impact inflation

One of the most common arguments that arise in conversations about metrics is that they can be easily 'gamed' or manipulated. Human nature being what it is, we have noticed a tendency to assign greater impact to a contribution than perhaps is warranted. In general, group leaders self-assess significant contributions and require input from the project leadership to claim an enabling impact. To avoid inflation, and ensure alignment between groups, the senior leaders of the CADD group meet quarterly to review and challenge the impact assignments for all significant or enabling contributions. In general, we find that a substantial number of contributions are downgraded at these meetings. These meetings are crucial to the success of the metrics process as we believe that they make the impact assessments less subjective. An additional tool for limiting the degree of inflation is that all categorizations are open-access, allowing any-



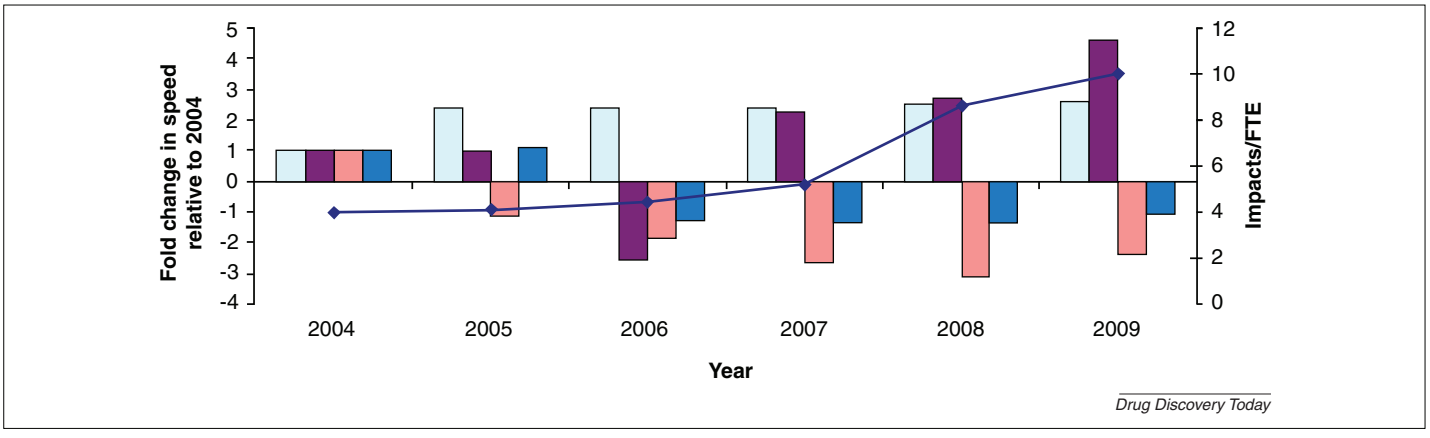


**FIGURE 6**

Effect of the software portfolio on impacts. The dollar amounts spent on each vendor are normalized to the corresponding amount spent in 2004. The graph shows the increases in financial commitments to vendors A (aqua bars) and B (purple bars) alongside an increasing rate of impacts/FTE (blue line). Over the same time period, dramatic reductions in commitments to all other vendors occurred (pink bars). In addition, the total software budget (blue bars) decreased relative to 2004 by up to 20% before returning to 2004 levels in 2009.

one within the greater technology organization to see the assignments for any contribution.

## Conclusion

Here, we have outlined our approach to metrics for an industrial CADD group. We have made a conscious choice to consider the quality of the impacts of our work over the quantity of tasks performed. One favorable aspect of this approach to metrics is that it provides freedom to the scientist to determine how best to address a particular question. The metrics themselves consider only the outcome of the work rather than the process to achieve it. We believe, and so far have observed, that such an approach to metrics will largely avoid the stifling effect on innovation that is often attributed to them [9].

Our approach to metrics collection is largely automated by leveraging pre-existing databases that senior management has already mandated. Although not effortless, the additional workload that assembling metrics has created has been more than compensated for by the information that we have gained on our overall effectiveness. Indeed, we have found the performance metrics to be valuable for allocating resources to projects, assessing management initiatives to streamline our work environment and understanding the differential value of innovative tools developed internally.

Attitudes toward metrics have shifted in our group since the process began. Although it would be unreasonable to claim that metrics are viewed warmly by the scientists in the group, the level of resistance to them has fallen dramatically since their initial implementation. Despite some continued skepticism, most group members view quantifiable performance metrics as a net benefit owing to how the metrics are being used to advocate for the group with senior management, improve staffing levels on projects, or allocate budget resources.

Finally, there are several additional extensions of these metrics that we are actively considering. For example, extending them to include information about the computational methodology used to generate a particular contribution. This would allow us to compare the effectiveness of various approaches to problems. However, it is important to consider the risk–reward trade-offs before engaging in finer iterations of measurement. The definition, collection and communication of each of these metrics consume resources that might be better allocated to scientific project work. Understanding the cost of a metric in the context of the value derived from it is crucial in deciding whether to pursue the development of new metrics.

## Acknowledgements

We would like to thank John Houston, Nicole Katz, Alastair Binnie, Martyn Banks and Stephen Brenner for many helpful discussions throughout this work.

## References

1 Munos, B. (2009) Lessons from 60 years of pharmaceutical innovation. *Nat. Rev. Drug Discov.* 8, 959–968
2 Zuckerman, D.S. (2006) *Pharmaceutical Metrics: Measuring and Improving R&D Performance.* Gower Publishing Company
3 Hwang, B.-G. *et al.* (2010) Performance metric development for pharmaceutical construction projects. *Int. J. Project Manag.* 28, 265–274
4 Kerr, S. (1975) On the folly of rewarding A, while hoping for B. *Acad. Manag. J.* 18, 769–783
5 Abbott, A.C. *et al.* (2010) Metrics: do metrics matter? *Nature* 465, 860–862
6 Van Noorden, R. (2010) Metrics: a profusion of measures. *Nature* 465, 864–866
7 Hirsch, J.E. (2005) An index to quantify an individual's scientific research output. *Proc. Natl. Acad. Sci. U.S.A.* 102, 16569–16572
8 Drucker, P.F. (1992) *Managing for the Future: the 1990s and Beyond.* Dutton
9 Ullman, F. and Boutellier, R. (2008) Drug discovery: are productivity metrics inhibiting motivation and creativity? *Drug Discov. Today* 13, 997–1001
10 Davenport, T.H. (2005) *Thinking for a Living: How to Get Better Performances and Results from Knowledge Workers.* Havard Business School Publishing
11 Hubbard, D.W. (2007) *How to Measure Anything: Finding the Value of 'Intangibles' in Business.* John Wiley & Sons
12 Houston John, G. *et al.* (2008) Case study: impact of technology investment on lead discovery at Bristol-Myers Squibb, 1998–2006. *Drug Discov. Today* 13, 44–51
13 Drucker, P.F. (1969) *The Age of Discontinuity.* Butterworth-Heinimann
14 Drucker, P.F. (1999) Knowledge-worker productivity: the biggest challenge. *Calif. Manag. Rev.* 41, 79–94

*Deborah Loughney,*
*Brian L. Claus,*
*Stephen R. Johnson*

*Computer-Assisted Drug Design, Molecular Sciences and Candidate Optimization, Bristol-Myers Squibb, PO Box 5400, Princeton, NJ 08543-5400, USA*