Improving Radiology Workflow with Automated Examination Tracking and Alerts

Oleg S. Pianykh, PhD[a,b], Christina Jaworsky[a], MT Shore[a], Daniel I. Rosenthal, MD[a,b]

Abstract

The modern radiology workflow is a production line where imaging examinations pass in sequence through many steps. In busy clinical environments, even a minor delay in any step can propagate through the system and significantly lengthen the examination process. This is particularly true for the tasks delegated to the human operators, who may be distracted or stressed. We have developed an application to track examinations through a critical part of the workflow, from the image-acquisition scanners to the PACS archive. Our application identifies outliers and actively alerts radiology managers about the need to resolve these problems as soon as they happen. In this study, we investigate how this real-time tracking and alerting affected the speed of examination delivery to the radiologist. We demonstrate that active alerting produced a 3-fold reduction of examination-to-PACS delays. Additionally, we discover an overall improvement in examination-to-PACS delivery, evidence that the tracking and alerts instill a culture where timely processing is essential. By providing supervisors with information about exactly where delays emerge in their workflow and alerting the correct staff to take action, applications like ours create more robust radiology workflow with predictable, timely outcomes.

Key Words: PACS, RIS, multivariate regression

J Am Coll Radiol 2017;14:937-943. Copyright © 2017 American College of Radiology

INTRODUCTION

Background

Delivering expeditious patient care requires efficiency from the time a patient enters the hospital to the time of discharge. In radiology, each examination must go through many sequential steps: scheduling the patient, beginning and completing image acquisition, sending the completed images from the scanner into the digital image archive, dictating the preliminary and final interpretation, delivering the reports, and submitting charges. Each of these steps depends on the completion of the previous; therefore, delays in any of them can impede the entire process.

[a]Department of Radiology, Massachusetts General Hospital, Boston, Massachusetts.

[b]Harvard Medical School, Boston, Massachusetts.

Corresponding author and reprints: Oleg S. Pianykh, PhD, Massachusetts General Hospital, 25 New Chardon Street, Office #470, Boston, MA 02114; e-mail: opianykh@mgh.harvard.edu.

The authors have no conflicts of interest related to the material discussed in this article.

Because of its sequential nature, the clinical and radiology workflow can be thought of as a complex production line [1], requiring on-time processing through many discrete steps. Viewing radiology under a manufacturing lens is not new, and health care institutions have already started implementing manufacturing-inspired systems [2]. For example, Odense University Hospital, one of the largest hospitals in Denmark, saw a marked increase in efficiency by implementing the Toyota Production System, also known as LEAN, a management strategy based on production lines [3]. The same "production line" approach is also supported by the theoretical analysis coming from disciplines such as the queueing theory [4,5]. As queueing theory demonstrates, when a service provider becomes busy and resource utilization runs high (a rather typical scenario for a crowded clinical facility), even a minor processing delay can result in a major increase of processing and waiting times [6]. Many of these delays cannot be avoided, or at least planned for. Therefore, the most practical approach is to identify and correct these delays as soon

1546-1440/17/$36.00 ■ http://dx.doi.org/10.1016/j.jacr.2017.03.019

as possible, before they propagate down the remaining workflow [7].

The Massachusetts General Hospital Department of Radiology performs around 90,000 examinations per month. Before our project described below, each step of this workflow was largely isolated; there was no centralized tracking and no workflow synchronization mechanism for the large volume of examinations. As a result, workflow delays were not caught in real time, and often were not addressed promptly.

To solve this problem, we developed an application named *Tempus Fugit* (or TF for short). TF tracks examinations, identifies outliers, and alerts radiology supervisors to take immediate actions. In this work, we studied the effect TF had on one of the most critical workflow metrics: *time-to-PACS* (TTP), defined as the amount of time it takes a technologist to send a completed examination from an acquisition device to the PACS digital archive. As in many clinical facilities, our completed examinations are not sent to PACS automatically (although this is technically possible); instead, they are first postprocessed and verified by the technologists. Therefore, submission to PACS becomes manual, resulting in the possibility of delayed or, worse, forgotten examinations.

We studied how TF examination tracking can be used to reduce the percentage of examinations with a delayed TTP time, and how TF's tracking system can be used to enforce timely examination processing, especially by alerting managers when action is necessary. Active alerting has already shown efficiency in health care [8], and bringing this idea to TF was expected to further reduce TTP for delayed examinations.

Methods

TF was developed as a department-wide time-tracking application. Running as a background server process, TF constantly monitors many of the most critical radiology workflow metrics such as image counts, patient wait times, completed but unread examinations, or examination TTP. TF identifies any examinations falling outside of expected metrics and reports them to its users. To accomplish this, TF queries real-time data from the departmental PACS, RIS, and reporting databases and performs on-the-fly calculations to find any current outliers. To provide sufficient information, a single TF query may process as many as 800,000 PACS and RIS records, merged together with nontrivial examination-tracking logic. In essence, TF functions as a real-time Big Data mining engine, combining large data volume and complex processing.

The first version of TF was released in spring 2015 as a website. Although this was a great step forward compared with the pre-TF "we just hope everything works well" era, the website approach was intrinsically *passive*: managers had to watch the site to see whether any new outliers were detected. Therefore, in late October 2015, *active* alerting by means of text pages was added for TTP longer than 40 minutes. The application initially alerted for CT examinations; the other modalities followed within a few months.

In the active (alerting) mode, when TF detects a new study with TTP exceeding 40 minutes, an initial alert is sent to the appropriate area manager's pager, including all the necessary information about the delayed examination (accession number, patient medical record number, and current TTP delay). If the examination is still not in PACS after an additional 40 minutes, the alert is escalated to the modality supervisor. If nothing happens again, TF will continue alerting managers and their supervisors every 40 minutes until the examination is in PACS.

The alert threshold of 40 minutes was chosen empirically and derived from the expected network speeds, examination postprocessing time, and manual steps required to send images. Taking these factors into account, we conservatively estimated that a vigilant technologist should take no longer than 40 minutes to send images to PACS, barring extraordinary circumstances. Numerically, the 40-minute threshold was approximately the 90th percentile of all examination-to-PACS transmissions, suggesting that using this value will not result in alert fatigue or a "crying wolf" effect [9-13].

To objectively measure the effectiveness of this approach, time points from more than 250,000 CT examinations between January 2014 and March 2016 were extracted from RIS and PACS records, and their historical TTP values and trends were analyzed. To best observe the change in the number of examinations with a prolonged TTP, we defined a *delayed TTP* as TTP greater than 50 minutes, thus allowing 10 extra minutes of action to take place after TF sends its first 40-minute alert.

The following sections present the results of our analysis. The historical study was done using Matlab and Python; the TF website and tracking engine were implemented in SQL and C# (ASP.NET).

ANALYSIS

Overall Impact

Figure 1 demonstrates how the percentage of PACS examinations with *delayed TTP* (TTP > 50 minutes)