Optimal Patient-specific Post-operative Surveillance for Vascular Surgeries

Peter I. Frazier, Shanshan Zhang, Pranav Hanagal
School of Operations Research & Information Engineering, Cornell University

Andrew J. Meltzer, Darren B. Schneider
Department of Surgery, Weill Cornell Medical College, Cornell University

This work was supported by the NSF, AFOSR, and a Cornell intercampus pilot grant

We consider loss of blood flow in the leg, due to narrowing of an artery (peripheral vascular disease)

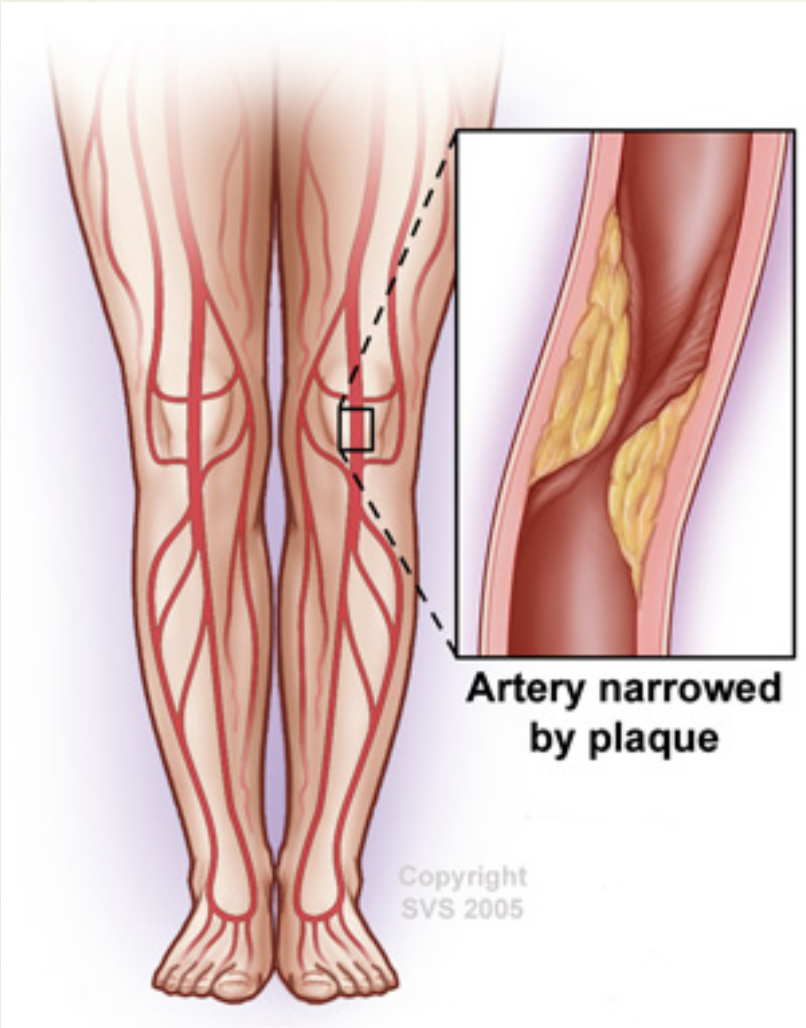

It causes pain and ulceration in the feet. It can require amputation, and can cause a coronary event.

Image credits: Society for vascular surgery; Wikipedia.

Surgery can fix this problem (but often, only temporarily)

- Surgical techniques include:
  - placing a stent to hold open the occluded artery
  - performing an angioplasty to open the artery

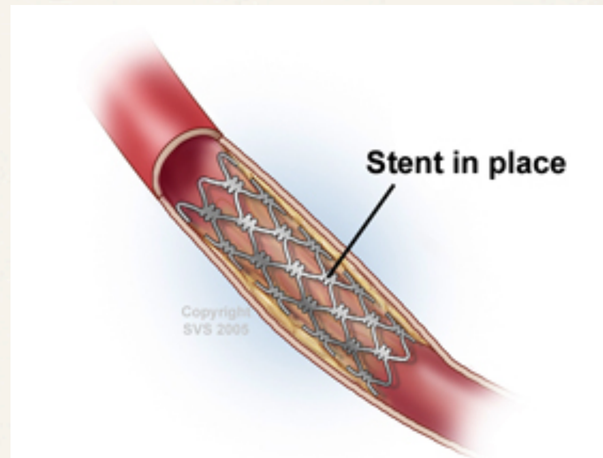

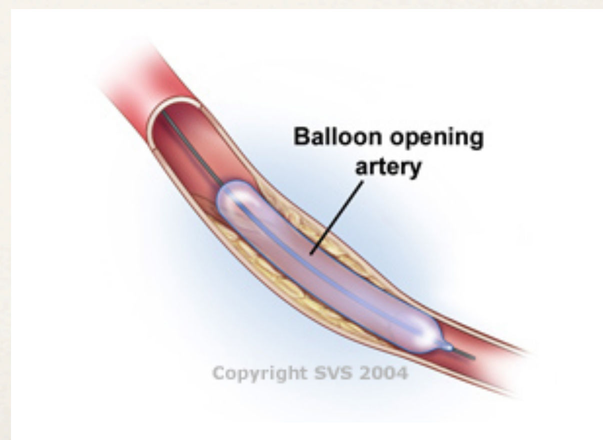

Later, the surgical repair may fail, causing blood to stop flowing

- Figure at right shows proportion of repairs that have not yet failed vs. time for two different patient groups (In this talk, we will focus on the Critical Limb Ischemia (CLI) group, and split it into two subgroups.)
- If this problem goes undetected for too long, the patient may suffer serious complications.
- Complications can require amputation of the affected limb, and sometimes result in death.

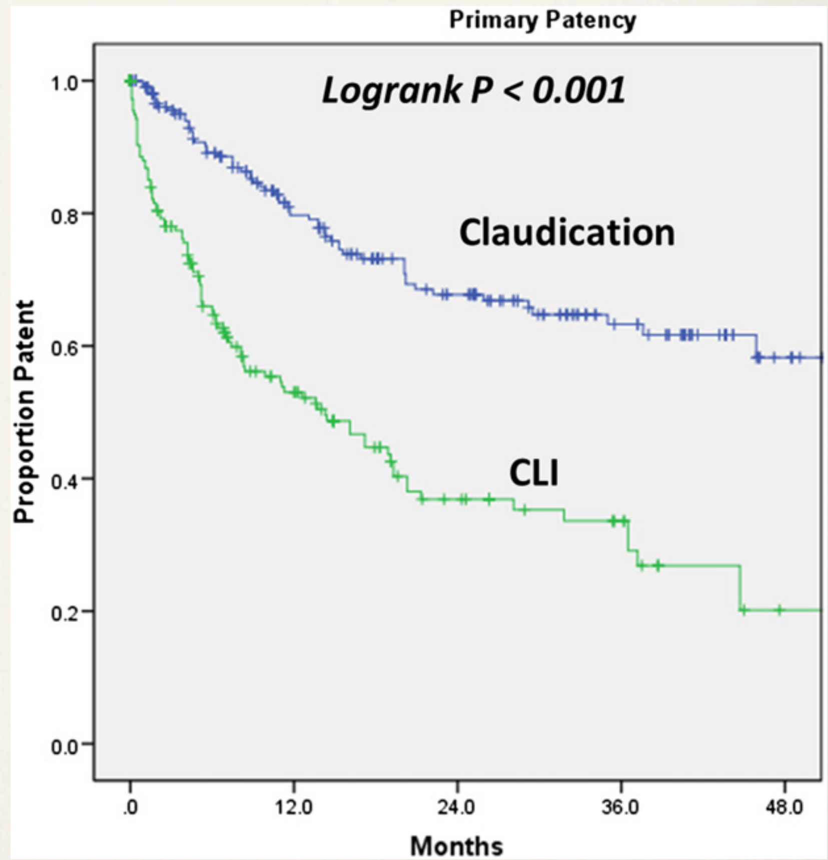

Meltzer et al., 2013

To catch failures before they cause problems, doctors perform surveillance

- Patients are tested after surgery to see whether the repair is still working (i.e., whether blood is still flowing).
- The test uses duplex ultrasound to compare the blood velocity in to and out of the leg.
- Tests are performed according to a fixed schedule, e.g., at months 1, 6, 12, 24, 36, 48, which is the same for all patients.

Current surveillance schedules can be improved

- Current practice is to use a one-size-fits-all approach:
  all patients get the same checkup (surveillance) schedule.
- For example, New York Presbyterian schedules checkups at 1 month, 6 months, 1 year , 2 years, 3 years, 4 years, ...
- Shouldn't more severe patients be checked more frequently?
- **Our goal**: develop optimal surveillance schedules, tuned to specific groups of patients.
    - We will consider patients with critical limb ischemia (CLI).
      Within this set of patients, we will consider two sub-groups:
      more severe disease (tissue loss); and less severe disease (rest pain).

Literature review

- There is a well-established literature on optimal inspection policies for deteriorating mechanical systems: [Barlow, Hunter & Proschan 1963], [Luss & Kander 1974], [Sengupta 1980], [Parmigiani 1994], [Parmigiani 1996], [Wang 2002], and a related and growing literature on inspection / surveillance / screening in medicine: at least 7 talks on screening and surveillance on Tuesday alone!
- Of these, the closest related work of which we are aware is [Sengupta 1980], which uses the same probabilistic model for how failures occur, but uses a different objective function.
- The statistical portion of our work is closely related to survival analysis, and censored regression.

We model the state of the repair with two random times

τ and δ are random times, whose distribution will be estimated from data, and will depend on patient type (to be discussed later).

Repair is working (blood is flowing)	Repair has failed, but is not symptomatic	Failure has caused emergency

0 Day of surgery

τ Repair fails

$\tau + \delta$ Failure causes emergency (possible amputation or death)

Time (t) →

We model the state of the repair with two random times

If a checkup occurs in the green region, no problem is detected.

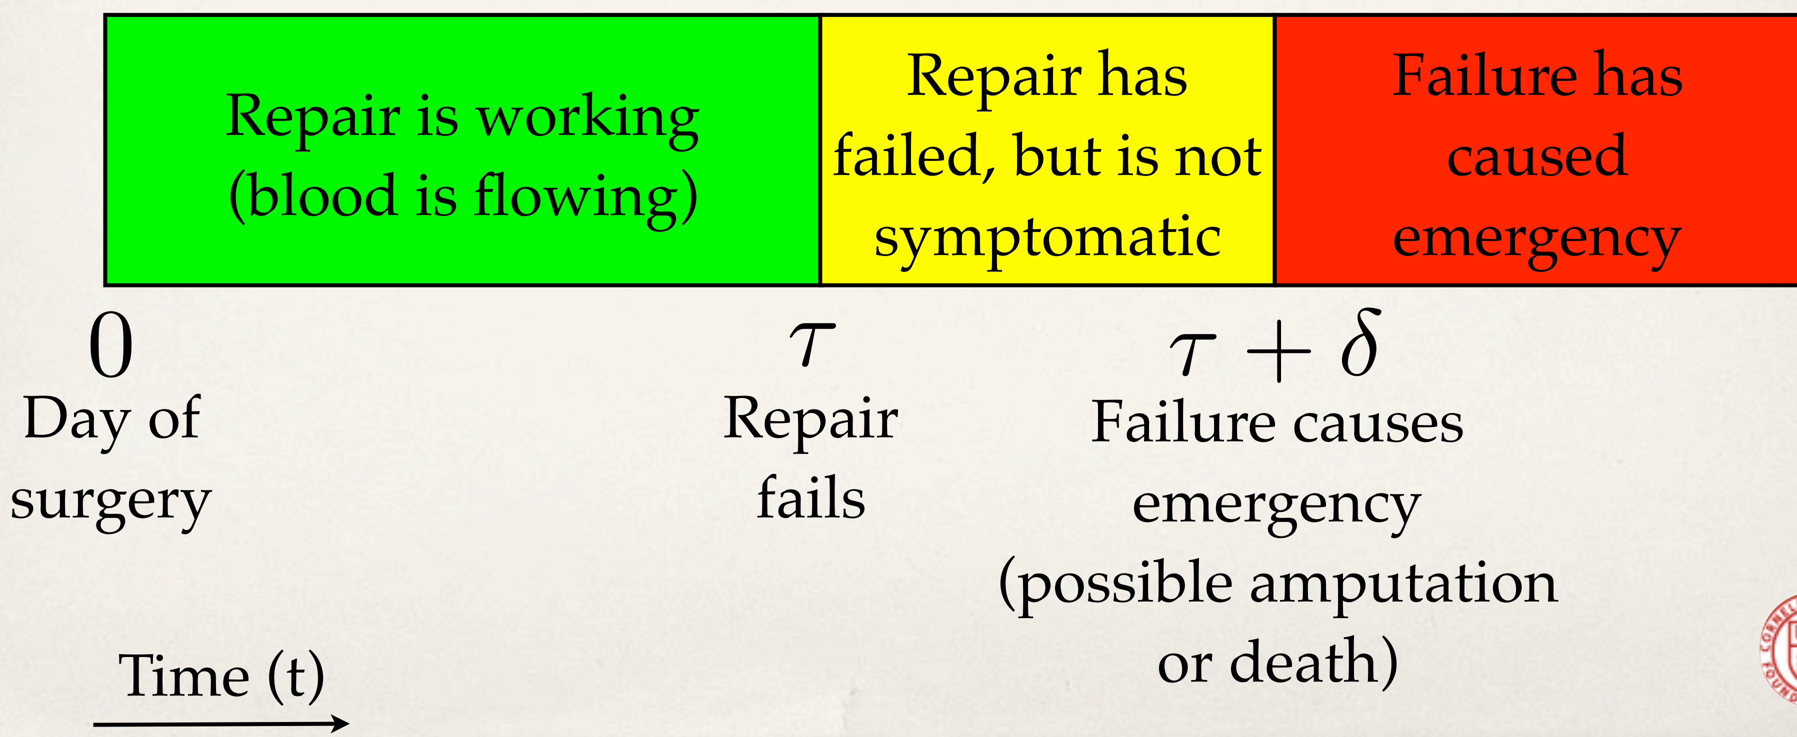

We model the state of the repair with two random times

If a checkup occurs in the yellow region, the failure of the repair is detected. Medical action can be taken in time, with relatively little risk to the patient.

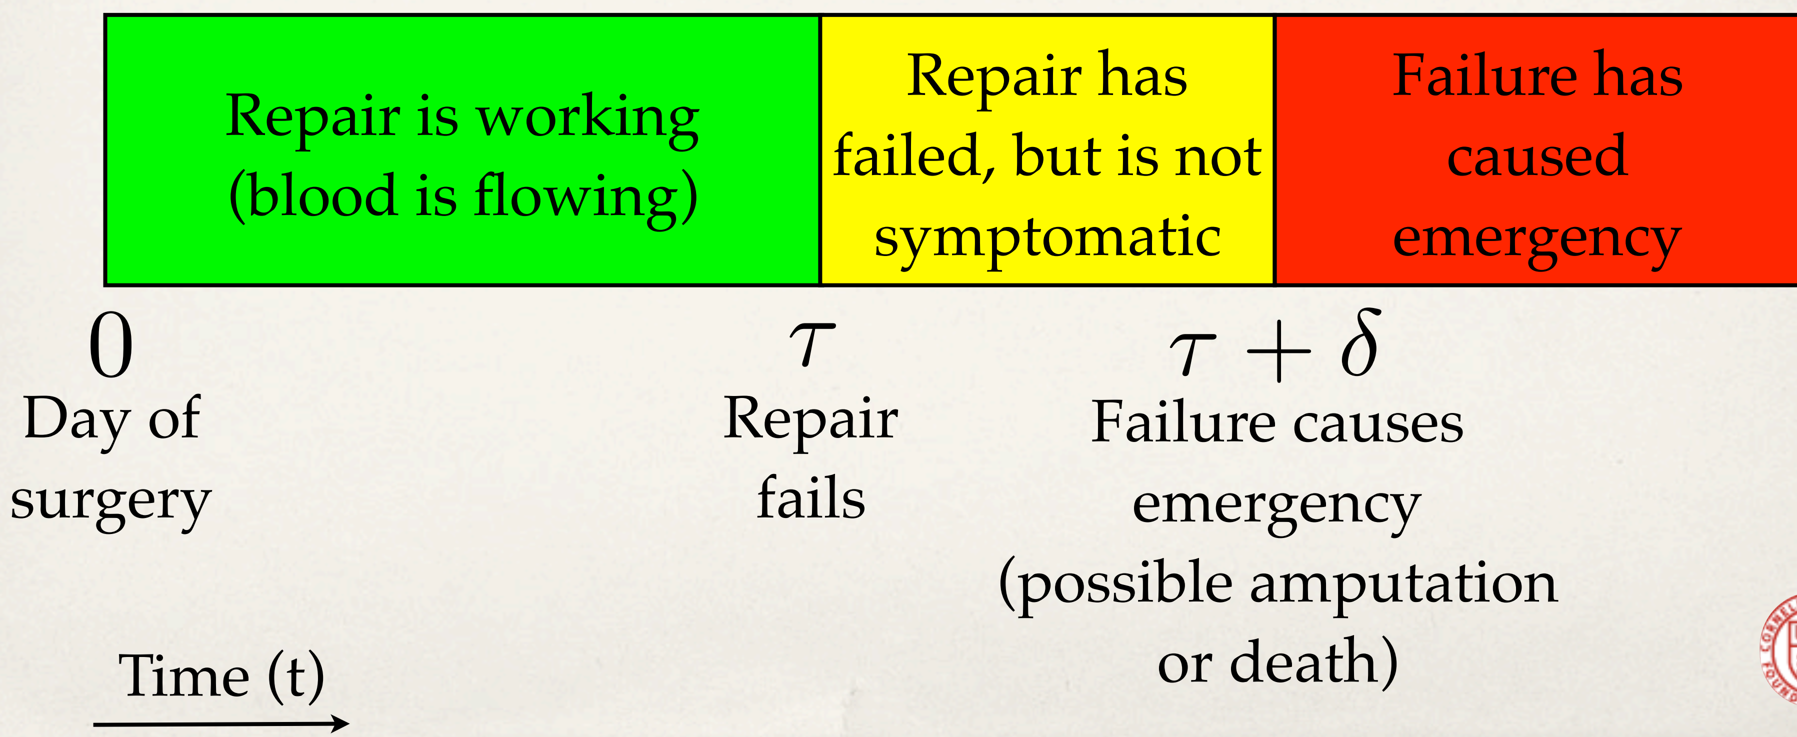

We model the state of the repair with two random times

If the failure of the repair goes undetected until $\tau + \delta$, the loss of blood flow becomes symptomatic, causing the patient to go to the emergency room with potentially serious complications. We wish to minimize the probability this occurs.

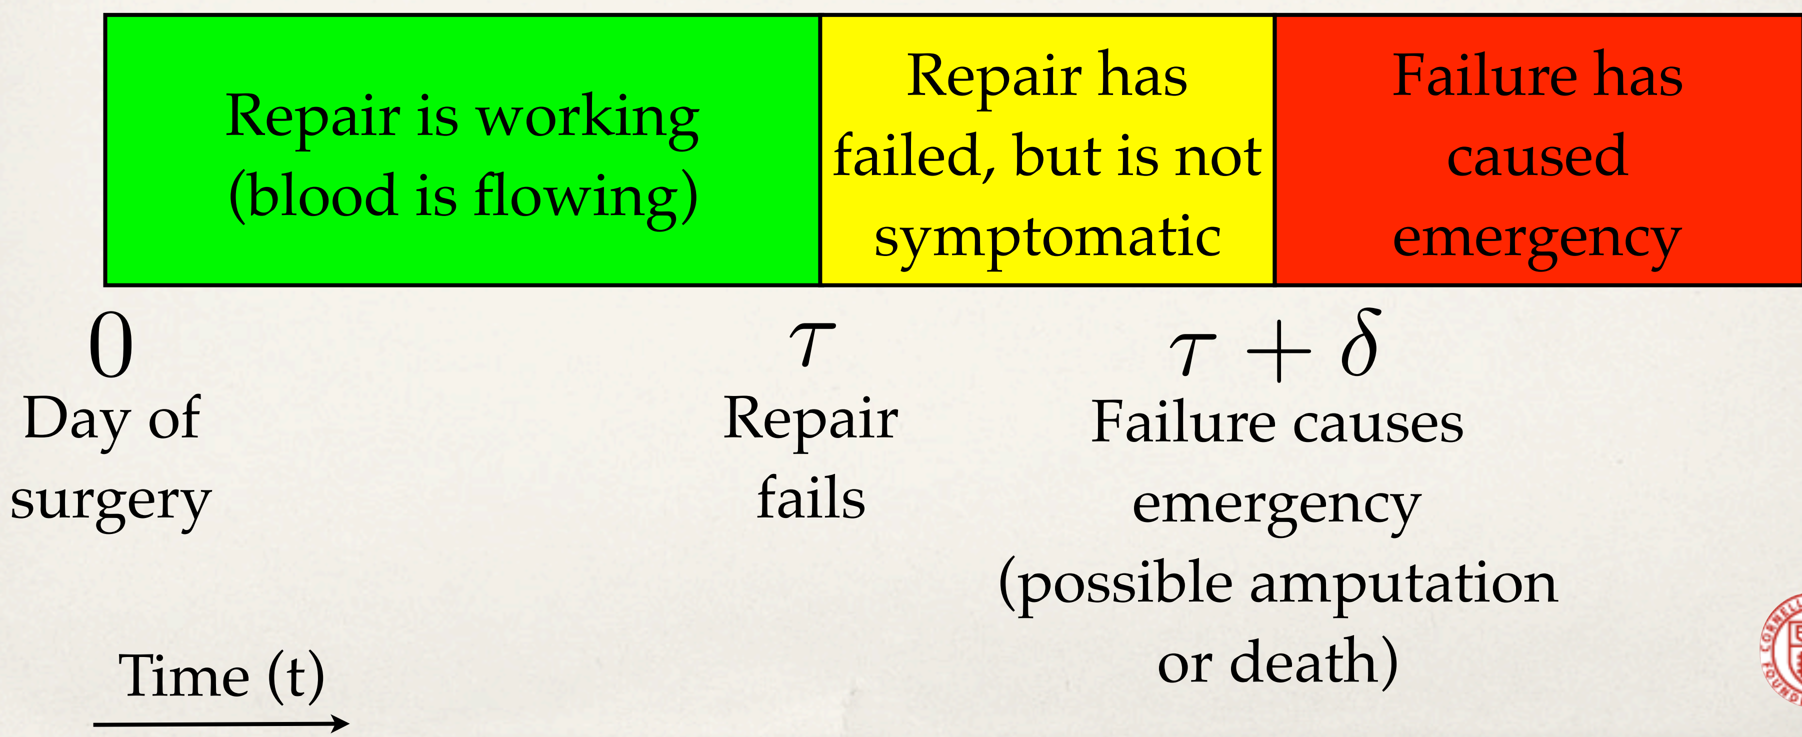

Our goal is to minimize the probability of an emergency

- Our goal is to choose the schedule so as to minimize the probability that an undetected failure causes an emergency to occur, up to a given time horizon T (24 months)

- We constrain the total number of checkups (typically 3 or 4).

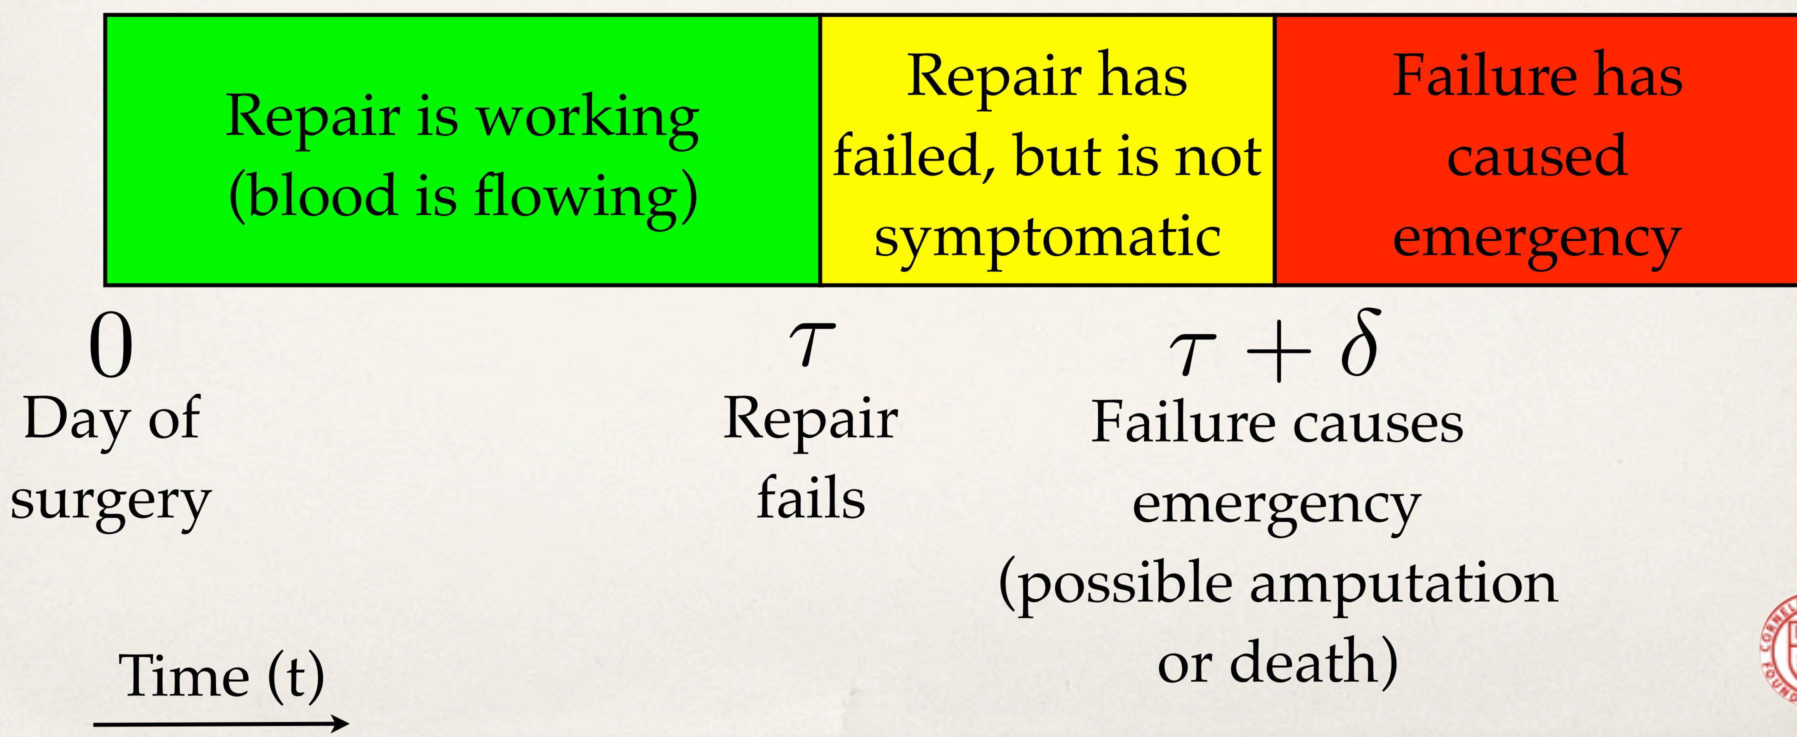

We provide an expression for the probability of an emergency

- Let our follow up schedule for a given patient be times t_1, t_2, ..., t_{m-1}, where m-1 is a constraint on the number of checkups allowed, let $T=t_m$ be our time horizon, and let $t_0=0$.

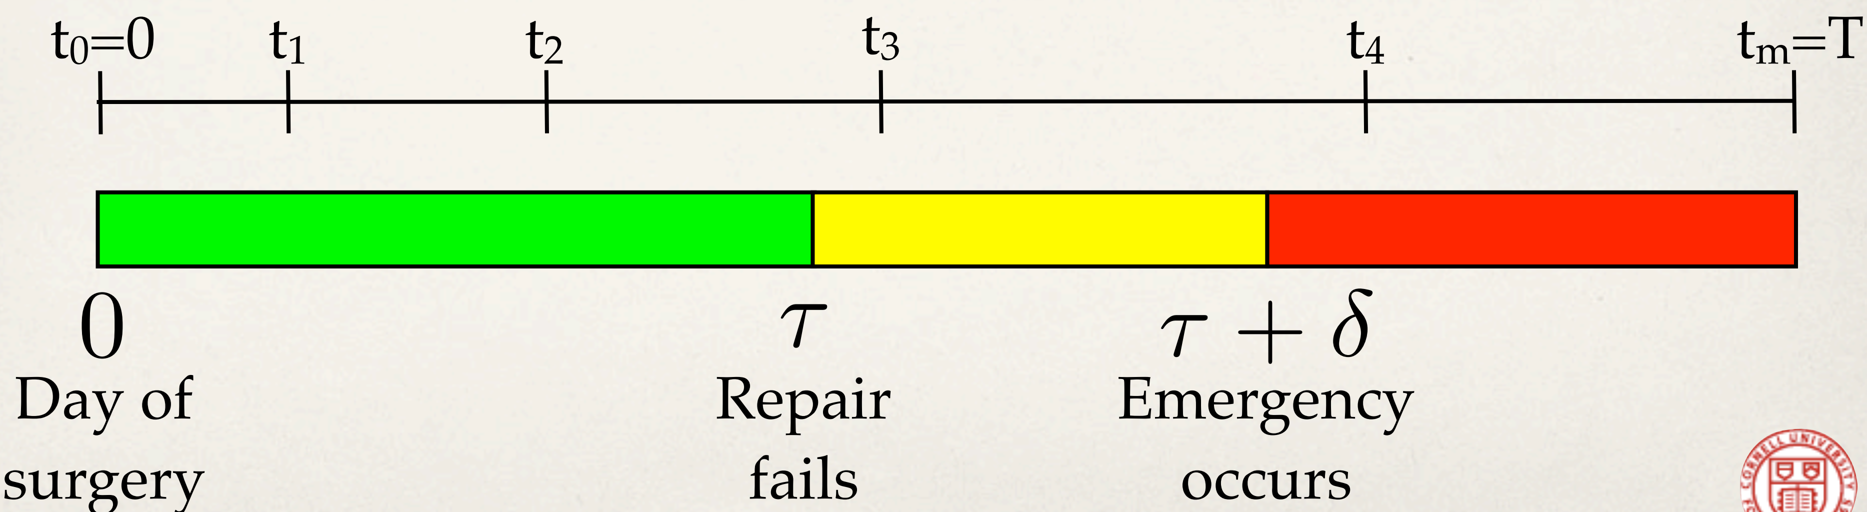

We provide an expression for the probability of an emergency

- An emergency occurs before time T if and only if $[\tau, \tau+\delta] \subseteq (t_{i-1}, t_i)$ occurs for some i.
- For the schedule below, for example, the first sample realization of τ and δ causes an emergency, while the second does not.

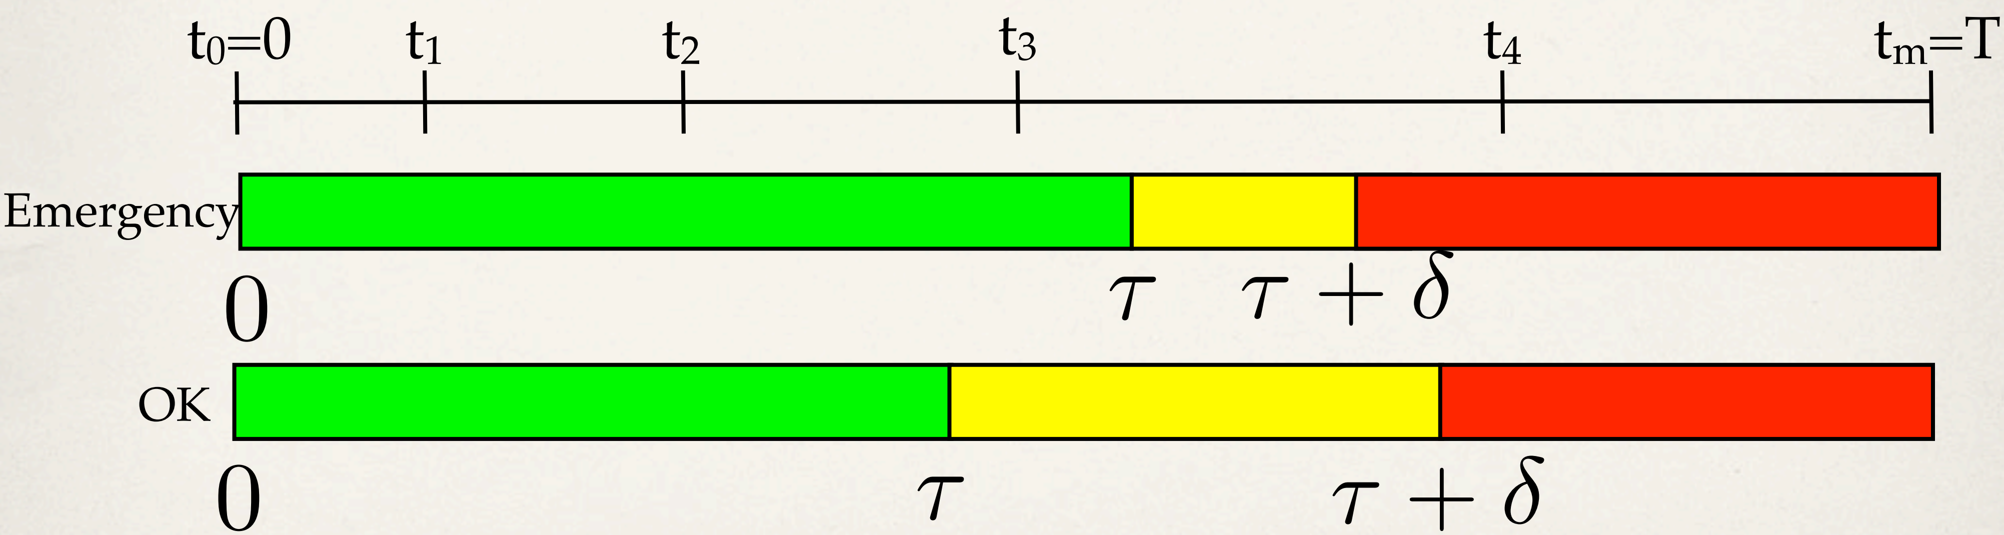

We provide an expression for the probability of an emergency

- We can rewrite the event $\{[\tau, \tau+\delta] \subseteq (t_{i-1}, t_i)\}$ as

$$\{\tau < t_{i-1}, \tau + \delta < t_i\} = \{t_{i-1} < \tau < t_i - \delta\}$$

- These events are disjoint across i, so the probability of an emergency is

$$\sum_{i=1}^{m} P\{t_{i-1} < \tau < t_i - \delta\}$$

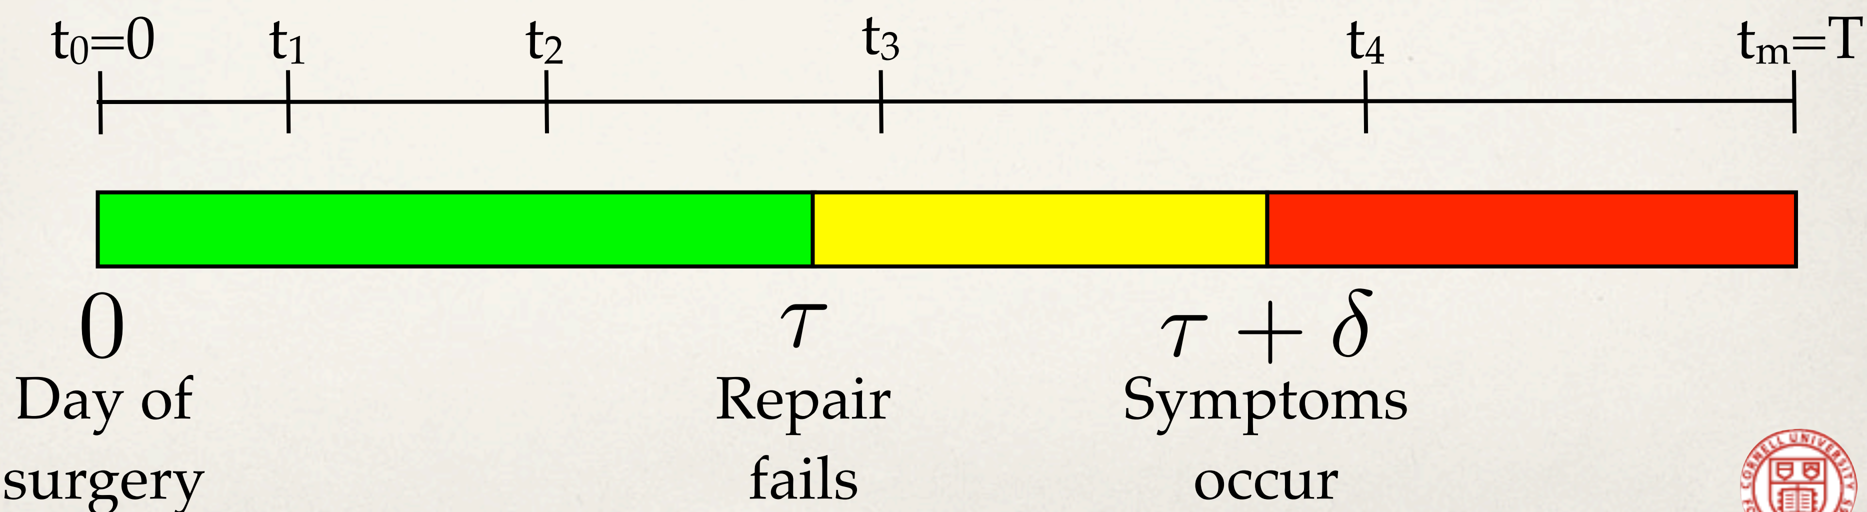

Minimizing the probability of an emergency corresponds to this optimization problem

$$
\begin{aligned}
\min_t F(t) = \sum_{i=1}^{m} & P(t_{i-1} < \tau < t_i - \delta) \\
\text{s.t.} \;\; & t_i < t_{i+1} \;\; \text{for} \;\; i = 0, \cdots, m-1, \\
& t_0 = 0, \\
& t_m = T.
\end{aligned}
$$

- Methodological questions:
  - Statistics: How can we estimate the distributions of τ and delta?
  - Optimization: How can we solve this optimization problem?

Overview of our statistical method

- We assume a parametric model, with τ independent of δ.

$$\tau \sim \mathrm{Gamma}(k, \lambda),$$
$$\delta \sim \mathrm{Exponential}(\alpha)$$

- We do not observe τ and δ directly, but instead censored versions of these variables

- We compute the likelihood corresponding to these censored observations, and then use maximum likelihood estimation to estimate tau and delta from data.

We have longitudinal data

- We have data on 239 patients from New York Presbyterian Hospital.
- For each patient, we have:
  - Duplex ultrasound date and result ("repair OK" or "repair failed"), for each duplex ultrasound performed. Note that patients deviate somewhat from the recommended schedule.
  - Emergency time ($\tau + \delta$), if an undetected failure caused one.
  - Censor date (e.g., date lost to follow up, or date of death for unrelated reasons), if no failures were detected.

We don't observe τ or δ directly; our observations are censored

- Each patient record falls into one of three cases: "No Failure", "Emergency Failure", or "Duplex-detected Failure"

We don't observe τ or δ directly; our observations are censored

- Here is the "**No Failure**" case (case 0):

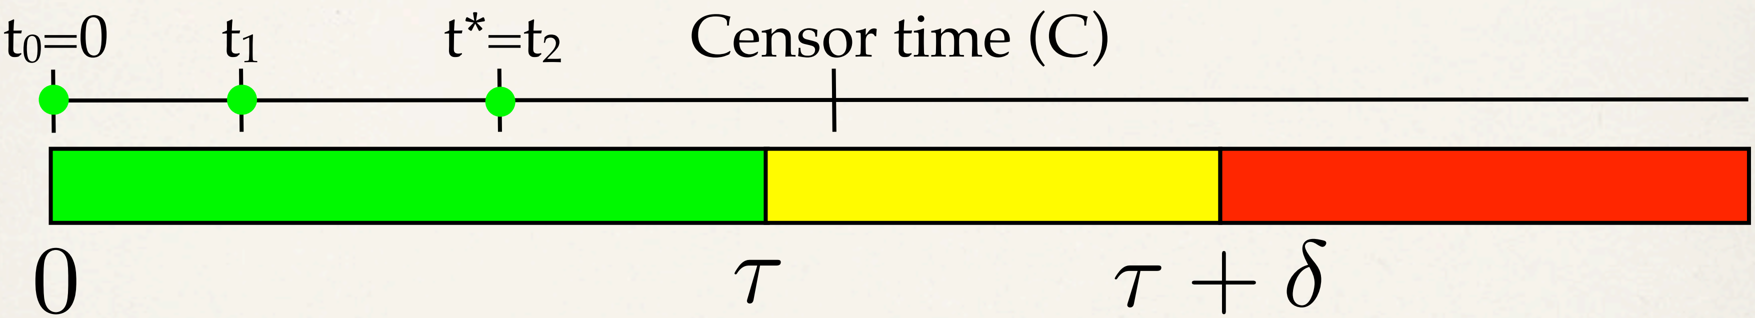

- We don't observe τ or $\tau + \delta$ for this patient. Instead, all we know is:

$$\tau \geq t^*, \tau + \delta \geq C$$

- The likelihood of the data for all patients in this case is:

$$\prod_{i \in \text{Case } 0} P\left\{\tau \geq t_i^*, \tau + \delta \geq C_i | \alpha, \lambda, k\right\}$$

We don't observe τ or δ directly; our observations are censored

- Here is the "**Emergency Failure**" case (case 1):

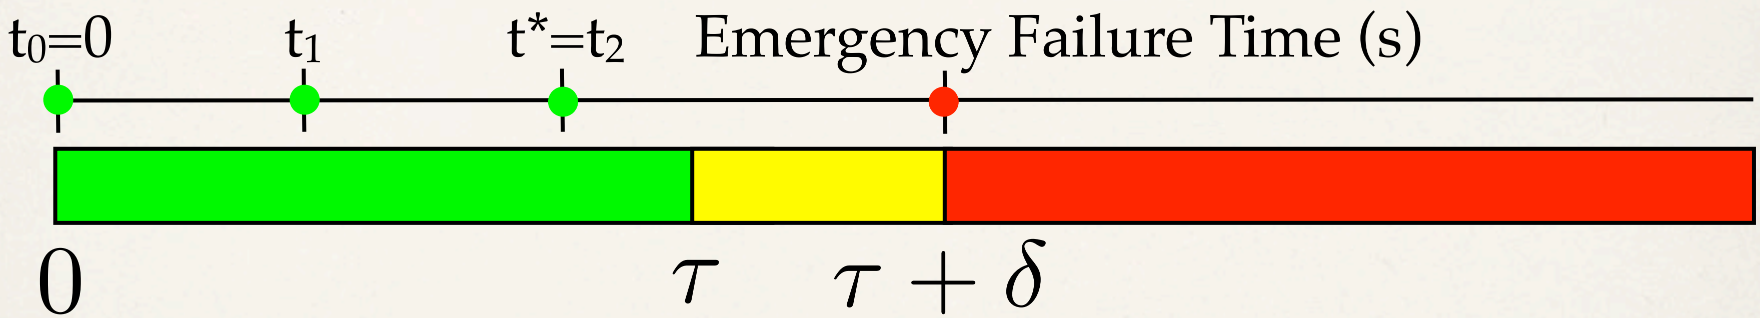

- We don't observe τ for this patient. Instead, all we know is:

$$\tau \geq t^*, \tau + \delta = s$$

- The likelihood of the data for all patients in this case is:

$$\prod_{i \in \text{Case } 1} P\left\{\tau \geq t_i^*, \tau + \delta = s_i | \alpha, \lambda, k\right\}$$

We don’t observe τ or δ directly; our observations are censored

- Here is the “**Failure detected early**” case (case 2):

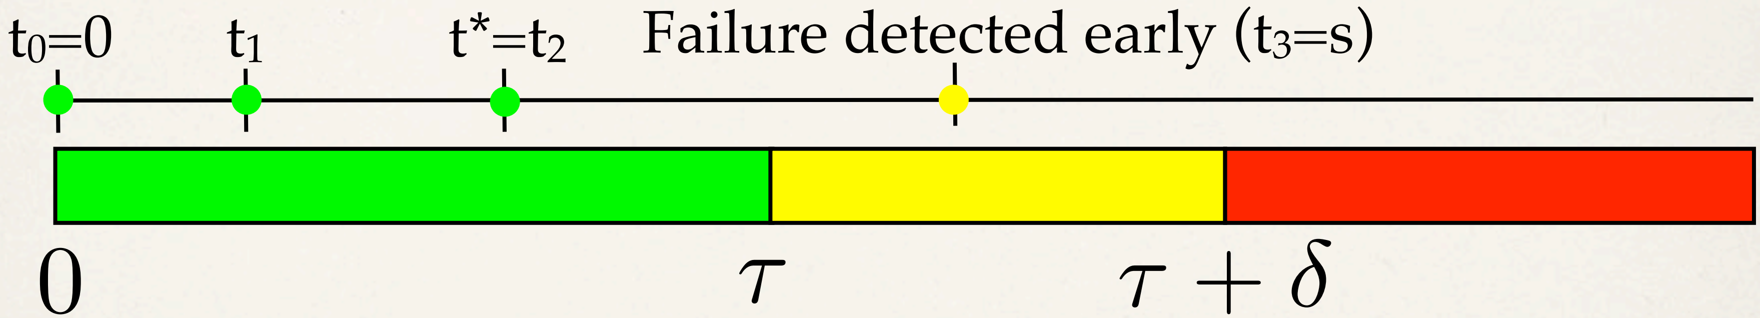

- We don’t observe τ or $\tau + \delta$ for this patient. Instead, all we know is:

$$t^* \leq \tau \leq s \leq \tau + \delta$$

- The likelihood of the data for all patients in this case is:

$$\prod_{i \in \text{Case } 2} P\left\{t_i^* \leq \tau \leq s_i \leq \tau + \delta | \alpha, \lambda, k\right\}$$

We use maximum likelihood estimation to estimate τ and δ

- Recall that $\tau \sim \text{Gamma}(k, \lambda),$
  $\delta \sim \text{Exponential}(\alpha)$
- We estimate k,λ, and α via

$$\arg\max_{k,\lambda,\alpha} \left[\prod_{i \in \text{Case } 0} P\left\{\tau \geq t_i^*, \tau + \delta \geq C_i | \alpha, \lambda, k\right\} \right] \times$$
$$\left[\prod_{i \in \text{Case } 1} P\left\{\tau \geq t_i^*, \tau + \delta = s_i | \alpha, \lambda, k\right\} \right] \times$$
$$\left[\prod_{i \in \text{Case } 2} P\left\{t_i^* \leq \tau \leq s_i \leq \tau + \delta | \alpha, \lambda, k\right\} \right]$$

- These probabilities can be calculated in closed form, and the optimization is done using fmincon with multiple restarts.

Overview of our optimization method

- Once we have estimated the parameters of our probability distribution, we plug them back into the optimization problem defining the optimal schedule:

$$\min_t F(t) = \sum_{i=1}^{m} P(t_{i-1} < \tau < t_i - \delta)$$

$$\text{s.t. } t_i < t_{i+1} \text{ for } i = 0, \cdots, m-1,$$

$$t_0 = 0, t_m = T.$$

- This is a continuous non-linear optimization problem with < 10 decision variables. The objective, and its 1st and 2nd derivatives can be computed analytically.
  - We can solve it using a multistart first or second order method.
  - For the special case k=1, we have developed a new method that is more efficient and guarantees a global optimum.

If τ is exponential, we can use a more efficient optimization method

- Proposition: if k=1 (so τ is exponential), then there is only one local minimizer (so it is the global minimizer), and it satisfies the following for i=1,2,...,m:

$$t_{i+1} = t_i - \frac{1}{\alpha} \log \left[1 - \frac{\alpha}{\alpha - \lambda} \left[1 - e^{-(\alpha-\lambda)(t_i - t_{i-1})} \right] \right]$$

If τ is exponential, we can use a more efficient optimization method

- Recall the formula from the previous slide:

$$t_{i+1} = t_i - \frac{1}{\alpha} \log \left[1 - \frac{\alpha}{\alpha - \lambda} \left[1 - e^{-(\alpha-\lambda)(t_i - t_{i-1})} \right] \right]$$

- Idea behind our algorithm:

  - If we were to guess at t_1, we could use the above expression to compute t_2 from (t_1, t_0), then t_3 from (t_2, t_1), etc., until we obtained t_m. One can show that this value of t_m is increasing in our guess at t_1.

  - At optimality, $t_m = T$. If we observed $t_m > T$ that means our guess at t_1 was too big. If we observed $t_m < T$, our guess at t_1 was too small.

  - Use this as a subroutine within a bisection algorithm to find the global optimum.

Using an optimal schedule decreases the probability of emergency

Patient Type	Schedule (in weeks)		Probability of Emergency (before T=2 years)	Improvement over status quo (%)
more severe (tissue loss)	status quo	4, 26, 52	0.51	-
	optimal	12, 26, 43	0.49	2.9%
	optimal+1	11, 23, 37, 56	0.46	9.4%
less severe (rest pain)	status quo	4, 26, 52	0.32	-
	optimal	20, 43, 69	0.28	9.7%
	optimal+1	16, 34, 52, 78	0.25	21.9%

Discussion

- The optimal schedules puts visits later for less severe patients.
- The status quo seems to be designed with more severe patients in mind. Less severe patients get the biggest advantage from switching to the optimal schedule.
- This analysis can also be used within a conversation about how many checkups to use for each patient type (Should we increase the number of follow-up visits for more severe patients?)

This analysis allows us to see what happens as we change the number of visits

More Severe Patients

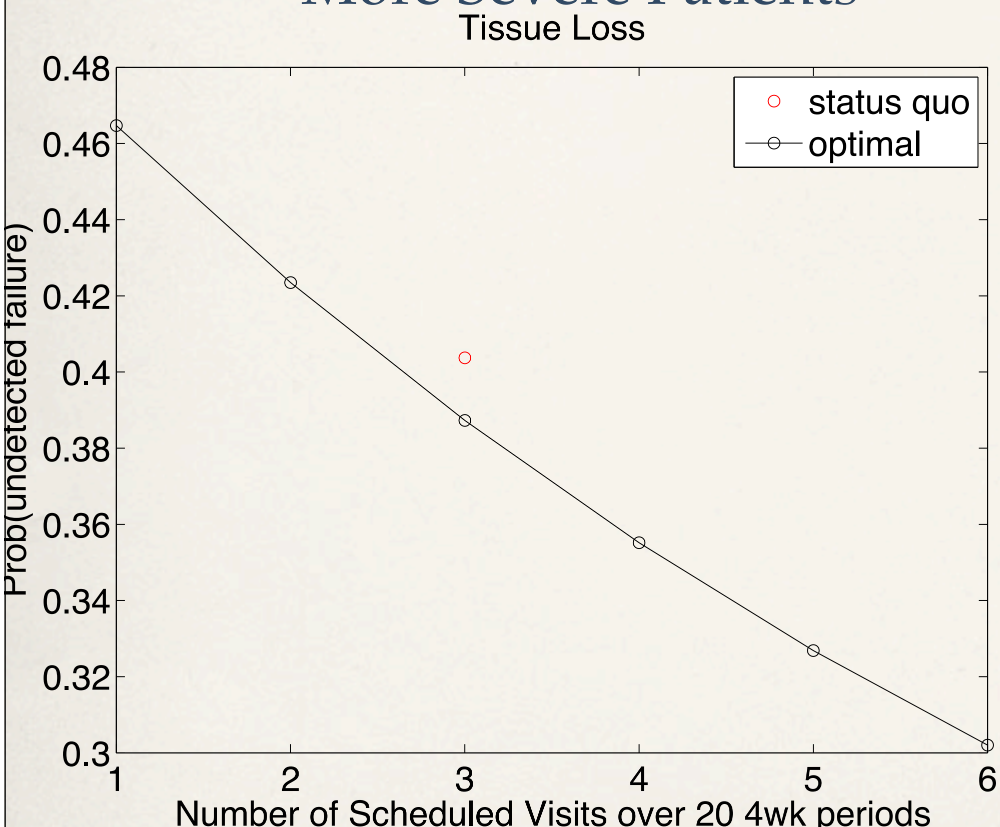

Less Severe Patients

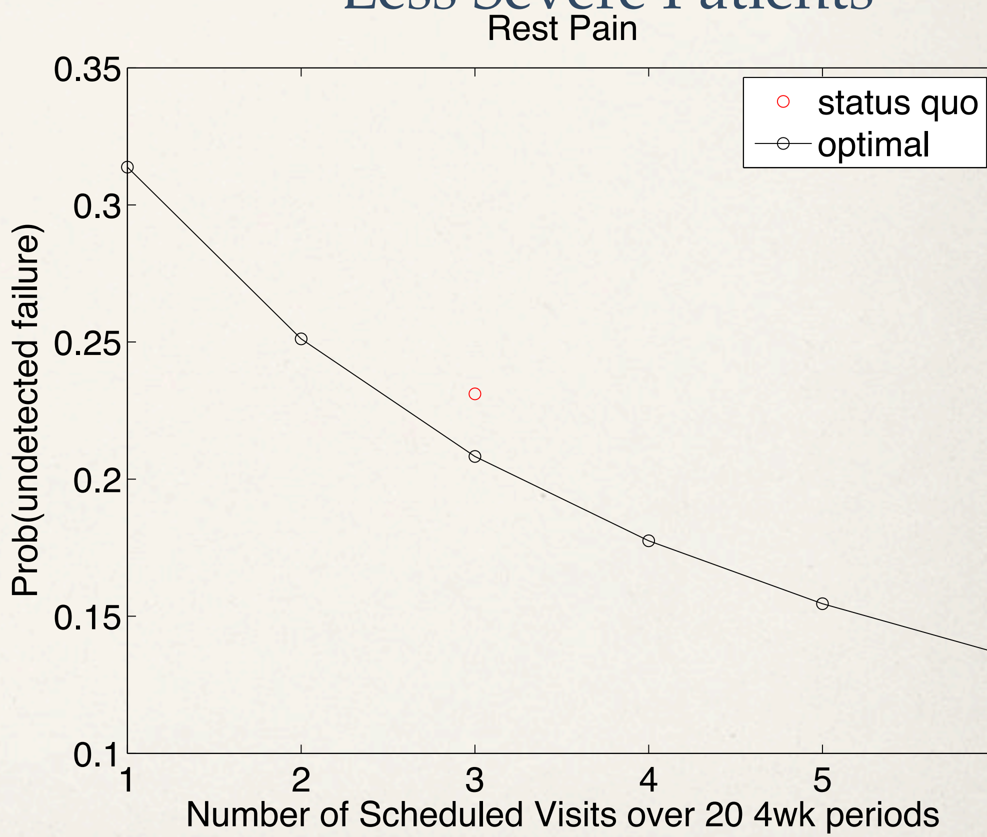

Note: these plots are for T=18 months

We are building an Excel spreadsheet to help doctors design schedules that are close to optimal and also convenient

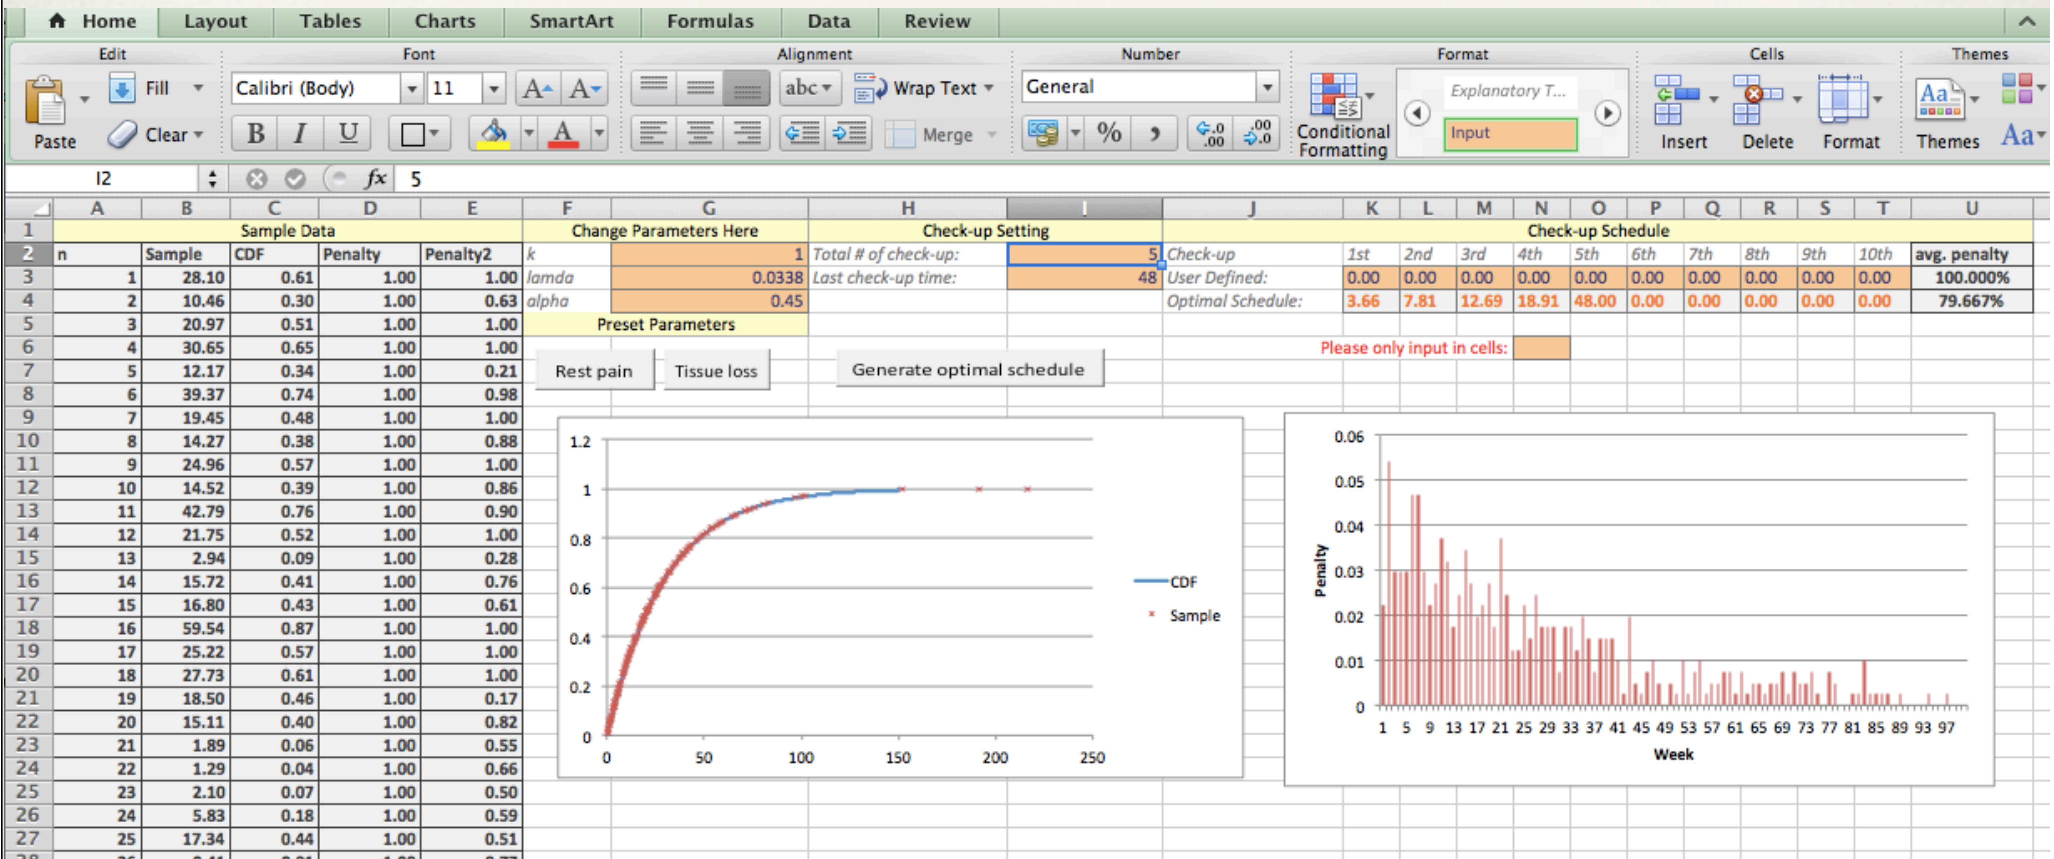

n	Sample	CDF	Penalty	Penalty2
1	28.10	0.61	1.00	1.00
2	10.46	0.30	1.00	0.63
3	20.97	0.51	1.00	1.00
4	30.65	0.65	1.00	1.00
5	12.17	0.34	1.00	0.21
6	39.37	0.74	1.00	0.98
7	19.45	0.48	1.00	1.00
8	14.27	0.38	1.00	0.88
9	24.96	0.57	1.00	1.00
10	14.52	0.39	1.00	0.86
11	42.79	0.76	1.00	0.90
12	21.75	0.52	1.00	1.00
13	2.94	0.09	1.00	0.28
14	15.72	0.41	1.00	0.76
15	16.80	0.43	1.00	0.61
16	59.54	0.87	1.00	1.00
17	25.22	0.57	1.00	1.00
18	27.73	0.61	1.00	1.00
19	18.50	0.46	1.00	0.17
20	15.11	0.40	1.00	0.82
21	1.89	0.06	1.00	0.55
22	1.29	0.04	1.00	0.66
23	2.10	0.07	1.00	0.50
24	5.83	0.18	1.00	0.59
25	17.34	0.44	1.00	0.51
26	0.41	0.01	1.00	0.77

Change Parameters Here		Check-up Setting	
k	1	Total # of check-up:	5
lamda	0.0338	Last check-up time:	48
alpha	0.45		
Preset Parameters			

Check-up	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th	avg. penalty
User Defined:	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.000%
Optimal Schedule:	3.66	7.81	12.69	18.91	48.00	0.00	0.00	0.00	0.00	0.00	79.667%

Future Work:

- **More flexible statistical models:** We would like to enrich to a richer class of parametric models that allow correlation between the random times, and use more patient covariates (especially co-morbidities).

- **More data**: We are in the process of obtaining a multi-institution dataset with more patient records.

- **Patient non-compliance**: Investigate the effect of patient non-compliance on the optimal schedule

Thank you!

- For more details, see:
  - S. Zhang, P. Hanagal, P.I. Frazier, A.J. Meltzer, D.B. Schneider, "Optimal Patient-specific Post- operative Surveillance for Vascular Surgery," 7th INFORMS Workshop on Data Mining and Health Informatics (DM-HI 2012), 2012.
  - A.J. Meltzer, A. Graham, P.H. Connolly, J.K. Karwowski, H.L. Bush, P.I. Frazier and D.B. Schneider, "Risk Factors for Early Failure after Peripheral Endovascular Intervention: Application of a Reliability Engineering Approach" Annals of Vascular Surgery, vol. 27, no. 1, pp 53–61, 2013.